Shallow Quantum Circuits with Uninitialized Ancillary Qubits

Yasuhiro Takahashi and Seiichiro Tani
NTT Communication Science Laboratories, NTT Corporation
{takahashi.yasuhiro,tani.seiichiro}@lab.ntt.co.jp

arXiv:1608.07020v2 [quant-ph] 16 Feb 2017

Abstract

We study the computational power of shallow quantum circuits with n input qubits, one output qubit, and two types of ancillary qubits: $O(\log n)$ initialized and $n^{O(1)}$ uninitialized qubits. The initial state of the uninitialized ancillary qubits is arbitrary, but we have to return their state into the initial one at the end of the computation. First, we show that such circuits can compute various symmetric functions on n bits, such as threshold functions. Then, we consider a polynomial-time probabilistic classical algorithm with an oracle that can perform such a circuit. We show that it can estimate the elements of any unitary matrix that can be implemented by a constant-depth quantum circuit on n qubits. Since it is not known whether these tasks can be done with only $O(\log n)$ initialized ancillary qubits, our results show the possibility that augmenting uninitialized ancillary qubits increases the computational power of shallow quantum circuits with only $O(\log n)$ initialized ancillary qubits. On the other hand, we give the cases where augmenting uninitialized ancillary qubits does not. More concretely, we consider near-logarithmic-depth quantum circuits with only $O(\log n)$ initialized ancillary qubits such that they include unbounded fan-out gates on a small number of qubits and unbounded Toffoli gates. We show that they cannot compute the parity function on n bits, even when they are augmented by $n^{O(1)}$ uninitialized ancillary qubits.

1 Introduction

1.1 Background and Main Results

Much attention has been paid to the computational power of shallow (i.e., polylogarithmic-depth) quantum circuits [5, 18, 12, 10, 13, 9, 3, 24, 23]. This line of research contributes not only to understanding the differences in computational power between such circuits and the corresponding classical ones, but also to implementing quantum algorithms on a quantum computer in the near future, which would be able to use qubits for only a short coherence time. When the computational power of shallow quantum circuits is discussed, it is usually assumed that they can have a large number of qubits initialized to $|0\rangle$ other than qubits for input data. This is because a large number of gates in shallow quantum circuits are applied in parallel and, in general, quantum circuits require a space for storing the computation history due to the laws of quantum mechanics. However, it seems difficult to prepare a large number of initialized ancillary qubits [8]. Thus, from a practical point of view, it is meaningful to study the computational power of shallow quantum circuits with only a small number of initialized ancillary qubits. It is also of theoretical interest to study the computational power of such circuits.

When shallow quantum circuits with n input qubits and one output qubit have only a small number of initialized ancillary qubits, their computational power seems quite low. For example, when the number of initialized ancillary qubits is $O(\log n)$, it is not known whether the OR function on n bits can be computed by a shallow quantum circuit consisting of one-qubit gates and CNOT gates, while it can be computed with depth $O(n)$ [1, 21]. Similarly, it is not known whether the threshold function on n bits can be computed by a shallow quantum circuit even with unbounded Toffoli gates and unbounded fan-out gates, while it can be computed with depth $n^{O(1)}$ [7]. It is natural to consider additional ancillary qubits satisfying the following two conditions: they should be easier to prepare than initialized ancillary qubits and have the possibility of increasing the computational power of shallow quantum circuits with only a small number of initialized ancillary qubits. As in [16], an interesting idea is to consider qubits in the completely mixed state as the additional ancillary qubits, but it is better not to assume any particular initial state in terms of preparation cost.

Very recently, Buhrman et al. [4] considered a classical computation with a small clean space and an additional space, which they call a catalytic-space computation, where the initial state of the

additional space is arbitrary, but its state has to be returned to the initial one at the end of the computation. They showed a surprising result: it appears that the catalytic log-space computation, which has a logarithmic-size clean space and a polynomial-size additional space, is more powerful than the standard log-space computation, which has only a logarithmic-size clean space. Motivated by this study and the above discussion, in this paper, we study the computational power of shallow quantum circuits with n input qubits, one output qubit, and two types of ancillary qubits: $O(\log n)$ initialized and $n^{O(1)}$ uninitialized qubits. The initial state of the uninitialized ancillary qubits is arbitrary, but we have to return their state into the initial one at the end of the computation.

Uninitialized ancillary qubits seem much easier to prepare than qubits with a particular initial state, i.e., they satisfy the above first condition on additional ancillary qubits. But do they satisfy the second condition? Specifically, do the $n^{O(1)}$ uninitialized ancillary qubits have the possibility of increasing the computational power of shallow quantum circuits with only $O(\log n)$ initialized ancillary qubits? They are known to be useful for constructing a few efficient quantum circuits [1, 22], but, to our knowledge, a complexity-theoretic analysis of quantum circuits with them has not yet been done.

First, to answer the above question affirmatively, we consider symmetric functions, which are Boolean functions whose output depends only on the number of ones in the input bits [15]. Let $\mathcal{E}$ be the class of symmetric functions f_n on n bits that are efficiently quantum computable from the number of ones in the input bits, i.e., $f_n \in \mathcal{E}$ if there exists a Boolean function g_m on $m = \lceil \log(n+1) \rceil$ bits such that $g_m(s) = f_n(x)$ for any $x \in \{0,1\}^n$, where $s \in \{0,1\}^m$ is the binary representation of the number of ones in x, and g_m can be computed by an $O(m)$-size quantum circuit consisting of the gates in $\mathcal{G}$. Here, $\mathcal{G}$ is the gate set consisting of a Hadamard gate, a $Z(\pm\pi/4)$ gate, a $Z'(\pm 2\pi/2^t)$ gate for any integer $t \geq 1$, and a CNOT gate, where $Z(\theta)$ is a phase-shift gate with angle $\theta \in \mathbb{R}$ and $Z'(\theta) = e^{-i\theta/2}Z(\theta)$. The class $\mathcal{E}$ includes various functions, such as the OR and threshold functions, for which it is not known whether there exist shallow quantum circuits with only $O(\log n)$ initialized ancillary qubits as described above. However, any $f_n \in \mathcal{E}$ can be computed by augmenting uninitialized ancillary qubits:

Theorem 1. *Any $f_n \in \mathcal{E}$ can be computed by an $O((\log n)^2)$-depth quantum circuit with n input qubits, one output qubit, $O(\log n)$ initialized ancillary qubits, and $O(n(\log n)^2)$ uninitialized ancillary qubits such that it consists of the gates in $\mathcal{G}$. Moreover, the depth decreases to $O(\log n)$ when the circuit is allowed to further include unbounded Toffoli gates on $O(\log n)$ qubits and unbounded fan-out gates.*

Theorem 1 shows the possibility that augmenting uninitialized ancillary qubits increases the computational power of shallow quantum circuits with only $O(\log n)$ initialized ancillary qubits. To show the usefulness of Theorem 1, we consider a shallow quantum circuit with $O(\log n)$ initialized and $n^{O(1)}$ uninitialized ancillary qubits such that it consists of the gates in $\mathcal{G}$ and threshold gates, where we assume that the number of threshold gates in each layer of the circuit is a constant. Such a circuit might be useful for solving some computational problems since threshold gates yield shallow quantum circuits for various arithmetic operations [25]. Suppose that we need to implement the circuit with only gates in $\mathcal{G}$. By replacing the threshold gates with the circuits for the threshold function in Theorem 1, we obtain a shallow quantum circuit with $O(\log n)$ initialized and $n^{O(1)}$ uninitialized ancillary qubits, while the above-mentioned previous circuit for the threshold function [7] yields an $n^{O(1)}$-depth one. The former is more implementable than the latter since the depth of the former is much smaller.

Then, to further show the strengths of uninitialized ancillary qubits, we consider a polynomial-time classical algorithm with an oracle that can perform a shallow quantum circuit with uninitialized ancillary qubits. When the algorithm sends a bit string w to the oracle, the oracle performs the circuit with input qubits initialized to $|w\rangle$ and sends back the classical outcome of the measurement on an output qubit. We deal with a matrix element estimation problem. Let $p(n)$ be a polynomial and C_n be an $O(1)$-depth quantum circuit on n qubits such that it consists of the gates in $\mathcal{G}$. The problem, denoted by $\mathrm{MEE}(p(n), C_n)$, is to compute $\alpha_x \in \mathbb{R}$ such that $|\alpha_x - |\langle 0^n|C_n|x\rangle|^2| \leq 1/p(n)$ for any input $x \in \{0,1\}^n$, where C_n also denotes the $2^n \times 2^n$ unitary matrix implemented by C_n. It is not known whether, for any $p(n)$ and C_n, $\mathrm{MEE}(p(n), C_n)$ can be solved by an $n^{O(1)}$-time classical algorithm (even in the bounded-error setting) [19]. This situation does not change when the algorithm has an oracle that can perform a shallow quantum circuit with only $O(\log n)$ initialized ancillary qubits. However, the problem can be solved by augmenting uninitialized ancillary qubits:

Theorem 2. *For any polynomial $p(n)$ and $O(1)$-depth quantum circuit C_n on n qubits such that it consists of the gates in $\mathcal{G}$, $\mathrm{MEE}(p(n), C_n)$ can be solved with probability exponentially (in n) close*

to 1 by an $n^{O(1)}$-time probabilistic classical algorithm with an oracle that can perform an $O(\log n)$-depth quantum circuit with $2n$ input qubits, one output qubit, no initialized ancillary qubits, and n uninitialized ancillary qubits such that it consists of the gates in $\mathcal{G}$.

As in Theorem 1, Theorem 2 shows the possibility that augmenting uninitialized ancillary qubits increases the computational power of shallow quantum circuits with only $O(\log n)$ initialized ancillary qubits. Here, we give a brief comment on the number of input qubits. If the number of input qubits is large, the algorithm can send 0^k for large k (besides another bit string) and the circuit can use a part of the input qubits as a large number of initialized ancillary qubits. To avoid this, when the length of the original input to the algorithm is n, we restrict the number of input qubits to $2n$.

Lastly, to show the limitations of uninitialized ancillary qubits, for any constant $0 \leq \delta < 1$, we consider $O((\log n)^\delta)$-depth quantum circuits consisting of one-qubit gates (not necessarily in $\mathcal{G}$), CNOT gates, unbounded fan-out gates on $(\log n)^{O(1)}$ qubits, and unbounded Toffoli gates. The proof of Bera [3] implies that such circuits with no ancillary qubits cannot compute the parity function on n bits. We extend this to the case with $O(\log n)$ initialized and $n^{O(1)}$ uninitialized ancillary qubits:

Theorem 3. *Let $0 \leq \delta < 1$ be an arbitrary constant. Then, the parity function on n bits cannot be computed by any $O((\log n)^\delta)$-depth quantum circuit with n input qubits, one output qubit, $O(\log n)$ initialized ancillary qubits, and $n^{O(1)}$ uninitialized ancillary qubits such that it consists of one-qubit gates,* CNOT *gates, unbounded fan-out gates on $(\log n)^{O(1)}$ qubits, and unbounded Toffoli gates.*

Theorem 3 gives the cases where augmenting uninitialized ancillary qubits does not increase the computational power of shallow quantum circuits. For example, the parity function cannot be computed by any $O((\log n)^{0.99})$-depth quantum circuit with only $O(\log n)$ initialized ancillary qubits, even when it is augmented by $n^{O(1)}$ uninitialized ancillary qubits. Moreover, Theorem 3 implies the following suggestion on how to improve Theorem 1. Theorem 1 yields $O(\log n)$-depth quantum circuits for various symmetric functions with $O(\log n)$ initialized and $n^{O(1)}$ uninitialized ancillary qubits such that they consist of the gates in $\mathcal{G}$, unbounded fan-out gates, and unbounded Toffoli gates. As described in Section 1.2, the circuits are based on a quantum circuit for computing the number of ones in the input bits, which can be considered as a quantum circuit for the parity function. Thus, they cannot be significantly improved simultaneously in terms of both the depth and the number of qubits on which unbounded fan-out gates act, since, otherwise, we would obtain a circuit that contradicts Theorem 3.

1.2 Overview of Techniques

A key ingredient of the circuits in Theorem 1 is a new shallow OR reduction circuit with $O(\log n)$ initialized and $n^{O(1)}$ uninitialized ancillary qubits. The circuit reduces the computation of the OR function on n bits to that on $m = \lceil \log(n+1) \rceil$ bits. Its first part is similar to the original OR reduction circuit with only a large number of initialized ancillary qubits [13] and yields a state that depends on the uninitialized ancillary qubits. Thus, in particular, we need to eliminate any dependency of the phase of the state on the uninitialized ancillary qubits. To do so, we add an appropriate phase to that of the state using a large number of uninitialized ancillary qubits. When we compute $f_n \in \mathcal{E}$ on input $x \in \{0,1\}^n$, we first compute the binary representation $s \in \{0,1\}^m$ of the number of ones in x by applying the quantum Fourier transform modulo 2^m to the output state of the OR reduction circuit, and then compute $g_m(s) = f_n(x)$, where g_m is the Boolean function associated with f_n described above.

We obtain the algorithm in Theorem 2 by modifying the $n^{O(1)}$-time probabilistic classical algorithm for $\mathrm{MEE}(p(n), C_n)$ with an oracle that can perform a commuting quantum circuit given in [19], where the circuit implements the Hadamard test [17]. A key ingredient of our algorithm is a shallow quantum circuit for the Hadamard test with uninitialized ancillary qubits. Although initialized ancillary qubits allow us to parallelize the Hadamard test [24], it has not been known whether uninitialized ancillary qubits are useful for this purpose. We show that uninitialized ancillary qubits can be used like initialized ones in parallelizing the Hadamard test. Replacing the commuting quantum circuit with the resulting shallow quantum circuit yields the desired algorithm.

We show Theorem 3 by extending the proof of Bera [3]. Our proof is different from the previous one in that it deals with ancillary qubits and unbounded fan-out gates. In particular, the key to Theorem 3 is to show that, for any quantum circuit C_n with $O(\log n)$ initialized and $n^{O(1)}$ uninitialized ancillary qubits such that it includes unbounded Toffoli gates, there exists an initial state of the uninitialized ancillary qubits such that C_n with the initial state is well approximated by $\tilde{C}_n$ with the same initial

state. Here, $\tilde{C}_n$ is the circuit obtained from C_n by removing unbounded Toffoli gates on a large number of qubits. Thus, roughly speaking, when C_n is a small-depth circuit for the parity function, so is $\tilde{C}_n$. This is impossible since $\tilde{C}_n$ does not have any gate on a large number of qubits and thus its output does not depend on all input qubits. The proof of Fang et al. [9] also shows the limitations of shallow quantum circuits, but it does not seem to work directly in our setting. This is because our circuits include unbounded fan-out gates and it seems difficult to approximate them with identity gates.

2 Preliminaries

2.1 Quantum Circuits and Uninitialized Ancillary Qubits

A quantum circuit consists of elementary gates, each of which is in an elementary gate set. The elementary gate set consists of one-qubit gates and a CNOT gate. Frequently used one-qubit gates in this paper are a Hadamard gate H, a phase-shift gate $Z(\theta)$, and its variant $Z'(\theta) = e^{-i\theta/2}Z(\theta)$. Here, $H = |+\rangle\langle 0| + |-\rangle\langle 1|$ and $Z(\theta) = |0\rangle\langle 0| + e^{i\theta}|1\rangle\langle 1|$ for any angle $\theta \in \mathbb{R}$, where $|\pm\rangle = (|0\rangle \pm |1\rangle)/\sqrt{2}$. In particular, we write $Z(\pi)$ as Z. Let $\mathcal{G}$ be the (universal) elementary gate set consisting of a H gate, a $Z(\pm\pi/4)$ gate, a $Z'(\pm 2\pi/2^t)$ gate for any integer $t \geq 1$, and a CNOT gate.

We also consider an elementary gate set including an unbounded fan-out gate and an unbounded Toffoli gate. Let $k \geq 1$ be an integer. A fan-out gate on $k+1$ qubits implements the operation defined as $|y\rangle \bigotimes_{j=1}^{k} |x_j\rangle \mapsto |y\rangle \bigotimes_{j=1}^{k} |x_j \oplus y\rangle$ for any $y, x_j \in \{0,1\}$, where $\oplus$ denotes addition modulo 2. The gate is a CNOT gate when $k = 1$. The first input qubit is called the control qubit. A k-controlled Toffoli gate implements the operation on $k+1$ qubits defined as $\left(\bigotimes_{j=1}^{k} |x_j\rangle\right) |y\rangle \mapsto \left(\bigotimes_{j=1}^{k} |x_j\rangle\right) |y \oplus \bigwedge_{j=1}^{k} x_j\rangle$ for any $y, x_j \in \{0,1\}$, where $\bigwedge$ denotes the logical AND. The gate is a CNOT gate when $k = 1$. The first k input qubits are called the control qubits and the last input qubit is called the target qubit. When it is permitted to apply a fan-out gate and a Toffoli gate on a non-constant number of qubits, they are called an unbounded fan-out gate and an unbounded Toffoli gate, respectively.

For convenience, we use a k-controlled $Z(\pm 2\pi/2^t)$ gate for any integer $t \geq 1$. The k-controlled $Z(2\pi/2^t)$ gate implements the operation on $k+1$ qubits defined as $\bigotimes_{j=1}^{k+1} |x_j\rangle \mapsto e^{\frac{2\pi i}{2^t} \bigwedge_{j=1}^{k+1} x_j} \bigotimes_{j=1}^{k+1} |x_j\rangle$ for any $x_j \in \{0,1\}$. We can choose an arbitrary qubit as the target qubit, and the other qubits are called the control qubits. The inverse of the gate is the k-controlled $Z(-2\pi/2^t)$ gate. It implements the same operation except that $2\pi i/2^t$ is replaced with $-2\pi i/2^t$. When it is permitted to apply these gates on a non-constant number of qubits, they are called unbounded $Z(\pm 2\pi/2^t)$ gates.

The complexity measures of a quantum circuit are its size and depth. The size of a quantum circuit is defined as the total size of all elementary gates in the circuit, where the size of an elementary gate is defined as the number of qubits on which the gate acts. To define the depth of a quantum circuit, we consider the circuit as a set of layers $1, \ldots, d$ consisting of elementary gates, where gates in the same layer act on pairwise disjoint sets of qubits and any gate in layer j is applied before any gate in layer $j+1$. The depth of the circuit is defined as the smallest possible value of d [10].

We deal with a uniform family of $n^{O(1)}$-size quantum circuits $\{C_n\}_{n\geq 1}$. The uniformity means that the function $1^n \mapsto \overline{C_n}$ can be computed by an $n^{O(1)}$-time classical algorithm, where $\overline{C_n}$ is the classical description of C_n. Each C_n has n input qubits and can have one output qubit and $n^{O(1)}$ ancillary qubits. The ancillary qubits are divided into two groups: $p = O(\log n)$ qubits and the remaining q qubits. We assume that, for any $x \in \{0,1\}^n$ and $y \in \{0,1\}$, we can initialize the input qubits and output qubit to $|x\rangle$ and $|y\rangle$, respectively. Moreover, we can initialize the p ancillary qubits to $|0\rangle$, which we call initialized ancillary qubits, but we cannot initialize the q ancillary qubits and do not know their initial state. The q ancillary qubits are called uninitialized ancillary qubits. When C_n has an output qubit, a measurement in the basis $\{|0\rangle, |1\rangle\}$ is performed on it at the end of the computation. The classical outcome of the measurement, which is 0 or 1, is called the output of C_n. A symbol denoting a quantum circuit also denotes its matrix representation in the computational basis. Any quantum circuit in this paper is understood to be an element of a uniform family of quantum circuits, and thus, for simplicity, we deal with a quantum circuit C_n in place of a family $\{C_n\}_{n\geq 1}$.

2.2 Computability of Boolean Functions

A Boolean function f_n on n bits is a mapping $f_n : \{0,1\}^n \to \{0,1\}$. The parity function on n bits, denoted by PA_n, is defined as $\mathrm{PA}_n(x) = 1$ if $|x|$ is odd and 0 otherwise, for any $x = x_1 \cdots x_n \in \{0,1\}^n$,

where $|x| = \sum_{j=1}^{n} x_j$. For any integer $1 \leq t \leq n$, the threshold function with a threshold t on n bits, denoted by TH_n^t, is defined as $\mathrm{TH}_n^t(x) = 1$ if $|x| \geq t$ and 0 otherwise. The OR function on n bits, denoted by OR_n, is TH_n^1. We define the computability of a Boolean function as follows:

Definition 1. *Let $f_n : \{0,1\}^n \to \{0,1\}$ be a Boolean function on n bits and C_n be a quantum circuit with n input qubits, one output qubit, p initialized ancillary qubits, and q uninitialized ancillary qubits. The circuit C_n computes f_n (or C_n is a quantum circuit for f_n) if the following conditions hold:*

- *For any $x \in \{0,1\}^n$ and $y \in \{0,1\}$, when the input qubits, output qubit, and p initialized ancillary qubits are initialized to $|x\rangle$, $|y\rangle$, and $|0^p\rangle$, respectively, the output of C_n is $y \oplus f_n(x)$ with probability 1, regardless of the initial state of the q uninitialized ancillary qubits.*
- *At the end of the computation, the state of all the qubits other than the output qubit is the same as their initial state.*

A Boolean function is called symmetric if its output depends only on the number of ones in the input bits [15]. We consider the class $\mathcal{E}$ of symmetric functions that are efficiently quantum computable from the number of ones in the input bits, where a classical version of the class is described in [14]. The precise definition of $\mathcal{E}$ is as follows:

Definition 2. *A symmetric function f_n on n bits is in $\mathcal{E}$ if there exists a Boolean function g_m on $m = \lceil \log(n+1) \rceil$ bits satisfying the following conditions:*

- *For any $x \in \{0,1\}^n$, $g_m(s) = f_n(x)$, where $s = s_1 \cdots s_m \in \{0,1\}^m$ such that $|x| = \sum_{k=1}^{m} s_k 2^{k-1}$.*
- *The function g_m can be computed by an $O(m)$-size quantum circuit with m input qubits, one output qubit, and $O(m)$ initialized ancillary qubits such that it consists of the gates in $\mathcal{G}$.*

It is easy to show that PA_n and TH_n^t (and thus OR_n) are in $\mathcal{E}$ since we can apply a NOT gate, i.e., $HZ(\pi/4)^4H$, and a 2-controlled Toffoli gate [1, 20].

The original OR reduction quantum circuit [13] is an $O(1)$-depth $O(n \log n)$-size one with n input qubits and $O(n \log n)$ initialized ancillary qubits such that it includes unbounded fan-out gates. The circuit reduces the problem of computing OR_n to that of computing OR_m, where $m = \lceil \log(n+1) \rceil$. We explain the idea of the circuit that will be used in our circuits. For any $x \in \{0,1\}^n$, let $|x\rangle$ be the initial state of the input qubits. The circuit transforms the state of the m initialized ancillary qubits into the state $\bigotimes_{k=1}^{m} |\varphi_k\rangle$ without changing the state of the other qubits. Here, $|\varphi_k\rangle = (|+\rangle + e^{\frac{2\pi i}{2^k}|x|}|-\rangle)/\sqrt{2}$ for any $1 \leq k \leq m$. If $|x| = 0$, $|\varphi_k\rangle = |0\rangle$ for any $1 \leq k \leq m$ and thus the output state is $|0^m\rangle$. If $|x| \geq 1$, $|\varphi_k\rangle = |1\rangle$ for some $1 \leq k \leq m$ and thus the output state is orthogonal to $|0^m\rangle$. The circuits with $n = 3$ are described in Section 3.1. Let $|x| = \sum_{k=1}^{m} s_k 2^{k-1}$ for some $s_k \in \{0,1\}$. It is easy to show that the state $\bigotimes_{k=1}^{m} |s_k\rangle$ can be obtained by applying $\mathrm{QFT}_{2^m}^{\dagger}$ to the state $\bigotimes_{k=1}^{m} H|\varphi_k\rangle$, where QFT_{2^m} is the quantum Fourier transform modulo 2^m and $\mathrm{QFT}_{2^m}^{\dagger}$ is its inverse.

3 Strengths of Uninitialized Ancillary Qubits

3.1 Our Idea for a New OR Reduction Circuit

A key ingredient of our circuits in Theorem 1 is an $O((\log n)^2)$-depth OR reduction circuit Q_n with $m = \lceil \log(n+1) \rceil$ initialized and $O(n(\log n)^2)$ uninitialized ancillary qubits such that it consists of the gates in $\mathcal{G}$. The depth decreases to $O(\log n)$ when the elementary gate set further includes an unbounded Toffoli gate on $O(\log n)$ qubits and an unbounded fan-out gate. The difficulty in constructing Q_n is how to utilize a large number of uninitialized ancillary qubits $B_1, \ldots, B_q$ to apply a large number of gates in parallel so that its output state is independent of the initial state of $B_1, \ldots, B_q$, where $q = n^{O(1)}$. The first part of Q_n is similar to the original OR reduction circuit [13] and yields a state that depends on the initial state of $B_1, \ldots, B_q$. We explain the idea for eliminating the dependency using the case where $n = 3$ (and thus $m = 2$) and the elementary gate set includes an unbounded fan-out gate and a 1-controlled $Z(2\pi/2^t)$ gate for $t = 1, 2$. Direct applications of the original OR reduction circuits for preparing $|\varphi_1\rangle$ and $|\varphi_2\rangle$ to the setting with uninitialized ancillary qubits are depicted in Figs. 1(a) and 1(b), respectively, where the number of initialized ancillary qubits is one and the other ancillary qubits, which are called B_1, B_2, B_3, are uninitialized ones.

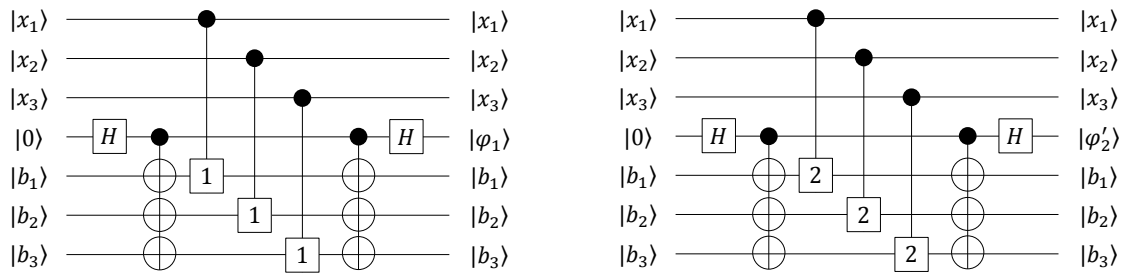

Figure 1: (a): Direct application of the original circuit for preparing $|\varphi_1\rangle$. The gate next to the H gate is a fan-out gate on four qubits, where the top qubit is the control qubit. For any integer $t \geq 1$, the gate represented as t is a $Z(2\pi/2^t)$ gate. (b): Direct application of the original circuit for $|\varphi_2\rangle$.

The direct application of the original circuit works for preparing $|\varphi_1\rangle$. In fact, when the initial state of B_1, B_2, B_3 is $|b_1\rangle|b_2\rangle|b_3\rangle$ for any $b_1, b_2, b_3 \in \{0,1\}$, a direct calculation shows that the circuit in Fig. 1(a) transforms the state of the initialized ancillary qubit into the state $(|+\rangle + e^{\frac{2\pi i}{2}\gamma}|-\rangle)/\sqrt{2}$, where we ignore the global phase and $\gamma = x_1 + x_2 + x_3 - 2(b_1x_1 + b_2x_2 + b_3x_3)$. Since $e^{\frac{2\pi i}{2}\gamma} = e^{\frac{2\pi i}{2}(x_1+x_2+x_3)}$, the obtained state is equal to $|\varphi_1\rangle$. The dependency on the initial state of B_1, B_2, B_3 is eliminated since the term $2(b_1x_1 + b_2x_2 + b_3x_3)$ yields only an angle of a multiple of 2π.

Unfortunately, the original circuit does not work for preparing $|\varphi_2\rangle$. A direct calculation shows that the circuit in Fig. 1(b) transforms the state of the initialized ancillary qubit into the state $(|+\rangle + e^{\frac{2\pi i}{2^2}\gamma}|-\rangle)/\sqrt{2}$, which is represented as $|\varphi'_2\rangle$ in Fig. 1(b). The phase $e^{\frac{2\pi i}{2^2}\gamma}$ depends on b_1, b_2, b_3. Our idea for eliminating the dependency is to use new uninitialized ancillary qubits and add an appropriate angle $\frac{2\pi}{2^2}\gamma'$ to $\frac{2\pi}{2^2}\gamma$, where $\gamma' = 2(b_1x_1 + b_2x_2 + b_3x_3) - 2^2(b'_1b_1x_1 + b'_2b_2x_2 + b'_3b_3x_3)$ and the initial state of the new uninitialized ancillary qubits is $|b'_1\rangle|b'_2\rangle|b'_3\rangle$ for any $b'_1, b'_2, b'_3 \in \{0,1\}$. We can obtain the angle $\frac{2\pi}{2^2}\gamma'$ using a quantum circuit similar to the one in Fig. 1(b) since γ' has a form similar to γ. Since $e^{\frac{2\pi i}{2^2}(\gamma+\gamma')} = e^{\frac{2\pi i}{2^2}(x_1+x_2+x_3)}$, the resulting state is equal to $|\varphi_2\rangle$ as desired. The dependency on the initial state of B_1, B_2, B_3 is eliminated since, as in the case of $|\varphi_1\rangle$, the term $2^2(b'_1b_1x_1 + b'_2b_2x_2 + b'_3b_3x_3)$ yields only an angle of a multiple of 2π.

We also use the following idea that contributes to decreasing the depth of Q_n. In the above explanation, we deal with two circuits for $|\varphi_1\rangle$ and $|\varphi_2\rangle$. We deal with m circuits in general and, to decrease the depth of Q_n, it is better to execute them in parallel. A simple method to do this is to first prepare m copies of the input state $|x_1\rangle \cdots |x_n\rangle$, but it is impossible to use this method since we have only $O(\log n)$ initialized ancillary qubits. Thus, we prepare the copies on uninitialized ancillary qubits, i.e., the states in the form $|x_1 \oplus a_1\rangle \cdots |x_n \oplus a_n\rangle$, where $|a_1\rangle \cdots |a_n\rangle$ is the initial state of the uninitialized ancillary qubits and $a_1, \ldots, a_n \in \{0,1\}$. To execute the m circuits in parallel, we use such copies on the uninitialized ancillary qubits and eliminate the dependency on their initial state.

3.2 Our OR Reduction Circuit and Proof of Theorem 1

We consider the case where the elementary gate set is $\mathcal{G}$ with an unbounded Toffoli gate on $O(\log n)$ qubits, an unbounded fan-out gate, and a k-controlled $Z(\pm 2\pi/2^t)$ gate for any integer $t \geq 1$. We prepare n input qubits $X_1, \ldots, X_n$ and m initialized ancillary qubits $I(1), \ldots, I(m)$. We also prepare $nm(m+3)/2$ uninitialized ancillary qubits, which are divided into two groups, A and B. Group A consists of mn qubits, which are divided into m groups $A(1), \ldots, A(m)$. Each $A(k)$ consists of n qubits $A_1(k), \ldots, A_n(k)$. Group B consists of $nm(m+1)/2$ qubits, which are divided into m groups $B(1), \ldots, B(m)$. Each $B(k)$ consists of kn qubits, which are divided into k groups $B(k,1), \ldots, B(k,k)$. Each $B(k,l)$ consists of n qubits $B_1(k,l), \ldots, B_n(k,l)$. For any $1 \leq k \leq m$, $|\varphi_k\rangle$ is prepared by using $I(k)$, $A(k)$, $B(k,l)$ for any $1 \leq l \leq k$ besides the input qubits.

The initial states of $X_1, \ldots, X_n$ are $|x_1\rangle, \cdots, |x_n\rangle$, respectively, for any $x_1, \ldots, x_n \in \{0,1\}$. For any $1 \leq k \leq m$, $I(k)$ is initialized to $|0\rangle$ and the initial states of $A_1(k), \ldots, A_n(k)$ are $|a_1(k)\rangle, \ldots, |a_n(k)\rangle$, respectively, for any $a_1(k), \ldots, a_n(k) \in \{0,1\}$. For any $1 \leq k \leq m$ and $1 \leq l \leq k$, the initial states of $B_1(k,l), \ldots, B_n(k,l)$ are $|b_1(k,l)\rangle, \ldots, |b_n(k,l)\rangle$, respectively, for any $b_1(k,l), \ldots, b_n(k,l) \in \{0,1\}$.

The circuit Q_n consists of m stages Stage $1, \ldots,$ Stage m, which are applied in this order. For any

$1 \le k \le m$, $|\varphi_k\rangle$ is prepared by using only Stage $1, \ldots,$ Stage k. As an example, Stages 1 and 2 with $n = 3$ are given in Appendix A.1. For any $1 \le s \le m$, Stage s is defined as follows:

1. Apply a H gate to $I(k)$ for every $s \le k \le m$ in parallel.
2. Apply a fan-out gate on $n+1$ qubits to $I(k)$ and $B_1(k,s), \ldots, B_n(k,s)$ for every $s \le k \le m$ in parallel, where $I(k)$ is the control qubit.
3. If $s = 1$, then do nothing. Otherwise, apply a fan-out gate on s qubits to $A_j(k)$ and $B_j(k,1), \ldots, B_j(k,s-1)$ for every $s \le k \le m$ and $1 \le j \le n$ in parallel, where $A_j(k)$ is the control qubit.
4. Apply an s-controlled $Z(-2\pi/2^{k-s+1})$ gate to $B_j(k,s)$ and the following qubits for every $s \le k \le m$ and $1 \le j \le n$ in parallel: If $s = 1$, then $A_j(k)$, where this qubit is the control qubit. Otherwise, $A_j(k)$ and $B_j(k,1), \ldots, B_j(k,s-1)$, where these qubits are the control qubits.
5. Apply a fan-out gate on $m-s+2$ qubits to X_j and $A_j(s), A_j(s+1), \ldots, A_j(m)$ for every $1 \le j \le n$ in parallel, where X_j is the control qubit.
6. Apply the inverse of the gates in Step 4.
7. Apply the gates in Step 5, Step 3, Step 2, and Step 1 (in this order).

The circuit Q_n outputs the desired state as follows. The proof can be found in Appendix A.1.

Lemma 1. *Let $x = x_1 \cdots x_n \in \{0,1\}^n$ be an input. For any $1 \le k \le m$ and $1 \le s \le k$, the state of $I(k)$ obtained by applying the gates from Stage 1 to Stage s is the state $(|+\rangle + e^{\frac{2\pi i}{2^k}\gamma(k,s)}|-\rangle)/\sqrt{2}$ and, in particular, the state of $I(k)$ is $|\varphi_k\rangle$ when $s = k$, where $\gamma(k,s) = |x| - 2^s \sum_{j=1}^{n} x_j \bigwedge_{l=1}^{s}(a_j(k) \oplus b_j(k,l))$. Moreover, for any qubit other than the initialized ancillary qubits, the state of the qubit obtained by applying those gates is the same as its initial state.*

By decomposing k-controlled $Z(\pm 2\pi/2^t)$ gates, we can regard Q_n as the circuit in the case where the elementary gate set is $\mathcal{G}$ with an unbounded Toffoli gate on $O(\log n)$ qubits and an unbounded fan-out gate [1]. However, the depth of the circuit is $O(\log n)$ times larger than that in Theorem 1. Thus, we replace such gates with their variants. A direct calculation shows that, for any integer $t \ge 1$ and $k \ge 1$, we can replace a k-controlled $Z(\pm 2\pi/2^t)$ gate in Q_n with a k-controlled $Z'(\pm 2\pi/2^t)$ gate without affecting the operation implemented by Q_n, where $Z'(\theta) = e^{-i\theta/2}Z(\theta)$. The k-controlled $Z'(2\pi/2^t)$ gate implements the operation on $k+1$ qubits defined as $\bigotimes_{j=1}^{k+1}|x_j\rangle \mapsto e^{-\frac{2\pi i}{2^{t+1}}(-1)^{x_{k+1}}\bigwedge_{j=1}^{k}x_j}\bigotimes_{j=1}^{k+1}|x_j\rangle$ for any $x_j \in \{0,1\}$. The last qubit is the target qubit, i.e., the qubit to which we apply a $Z'(2\pi/2^t)$ gate, and the other qubits are called the control qubits. The inverse of the gate is the k-controlled $Z'(-2\pi/2^t)$ gate and implements the same operation except that $-2\pi i/2^{t+1}$ is replaced with $2\pi i/2^{t+1}$.

The circuit Q_n has the desired complexity as follows. The proof can be found in Appendix A.2.

Lemma 2. *The depth of Q_n is $O((\log n)^2)$ when the elementary gate set is $\mathcal{G}$. It decreases to $O(\log n)$ when the elementary gate set further includes an unbounded Toffoli gate on $O(\log n)$ qubits and an unbounded fan-out gate. The circuit Q_n uses $O(\log n)$ initialized and $O(n(\log n)^2)$ uninitialized ancillary qubits in both of the cases.*

When we compute $f_n \in \mathcal{E}$ on input $x \in \{0,1\}^n$, we first compute the binary representation $s \in \{0,1\}^m$ of $|x|$. This can be done by applying $\mathrm{QFT}_{2^m}^{\dagger}$ to the state $\bigotimes_{k=1}^{m} H|\varphi_k\rangle$, which is simply obtained from the output state of Q_n. We then compute $g_m(s) = f_n(x)$ using only $O(\log n)$ initialized ancillary qubits, where g_m is the Boolean function associated with f_n (see Definition 2). This algorithm with Lemmas 1 and 2 implies Theorem 1. The details can be found in Appendix A.3.

3.3 Proof of Theorem 2

Let $p(n)$ be a polynomial and C_n be an $O(1)$-depth quantum circuit on n qubits such that it consists of the gates in $\mathcal{G}$. As described in Section 1, the problem $\mathrm{MEE}(p(n), C_n)$ is to compute $\alpha_x \in \mathbb{R}$ such that $|\alpha_x - |\langle 0^n|C_n|x\rangle|^2| \le 1/p(n)$ for any input $x \in \{0,1\}^n$, where C_n also denotes the $2^n \times 2^n$ unitary matrix implemented by C_n. For any $w = w_1 \cdots w_n \in \{0,1\}^n$, we define $Z(w) = \bigotimes_{j=1}^{n} Z_j^{w_j}$, where Z_j is applied

on the j-th qubit of C_n. Moreover, for any $x, w \in \{0,1\}^n$, we define $F_n(x,w) = \langle x|C_n^\dagger Z(w) C_n|x\rangle$. As shown in [19], MEE$(p(n), C_n)$ can be solved with probability exponentially (in n) close to 1 if there exists a probabilistic algorithm A_{F_n} such that, for any $x, w \in \{0,1\}^n$, the probability that $|A_{F_n}(x,w) - F_n(x,w)| \leq 0.5/p(n)$ is exponentially close to 1. In fact, due to the Chernoff-Hoeffding bound, the probabilistic algorithm for MEE$(p(n), C_n)$ is described with some $K = n^{O(1)}$ as follows, where the input is $x \in \{0,1\}^n$: For every $1 \leq j \leq K$, choose $w(j) \in \{0,1\}^n$ uniformly at random and compute $A_{F_n}(x, w(j))$. After that, output $(1/K)\sum_{j=1}^{K} A_{F_n}(x, w(j))$.

The probabilistic algorithm A_{F_n} in [19] can be considered as a repetition of a commuting quantum circuit D_{2n} for the Hadamard test with $2n$ input qubits and one output qubit. For any $x, w \in \{0,1\}^n$, the output of D_{2n} with the input qubits initialized to $|x\rangle|w\rangle$ and output qubit initialized to $|0\rangle$ is 0 with probability $(1 + F_n(x,w))/2$ (and thus 1 with probability $(1 - F_n(x,w))/2$). Thus, when the outputs 0 and 1 are regarded as 1 and -1, respectively, due to the Chernoff-Hoeffding bound, A_{F_n} is described with some $L = n^{O(1)}$ as follows, where the input is the pair of x and w, which are in $\{0,1\}^n$: For every $1 \leq j \leq L$, perform D_{2n} with the input qubits initialized to $|x\rangle|w\rangle$ and output qubit initialized to $|0\rangle$, and obtain its output $z_j(x,w) \in \{1,-1\}$. After that, output $(1/L)\sum_{j=1}^{L} z_j(x,w)$.

Our idea for proving Theorem 2 is to construct a parallelized version of the Hadamard test, denoted by E_{2n}, by using uninitialized ancillary qubits and replace D_{2n} in the above probabilistic algorithm for MEE$(p(n), C_n)$ with E_{2n}. Although the standard Hadamard test is a sequential application of controlled gates with the same control qubit, roughly speaking, E_{2n} first prepares the copies of the state of the control qubit on uninitialized ancillary qubits and then applies the controlled gates in parallel by using the copies. To describe E_{2n} more precisely, let $x = x_1 \cdots x_n, w = w_1 \cdots w_n \in \{0,1\}^n$. We prepare $2n$ input qubits $X_1, \ldots, X_n, W_1, \ldots, W_n$, one output qubit Y, and n uninitialized ancillary qubits $G(1), \ldots, G(n)$. The initial states of $X_1, \ldots, X_n, W_1, \ldots, W_n$ are $|x_1\rangle, \ldots, |x_n\rangle, |w_1\rangle, \ldots, |w_n\rangle$, respectively. The initial state of Y is $|0\rangle$ and the initial state of the uninitialized ancillary qubits is arbitrary. As an example, E_{2n} with $n = 3$ is given in Appendix A.4. The circuit E_{2n} is defined as follows, where we use fan-out gates and 2-controlled Z gates for simplicity:

1. Apply a H gate to Y.
2. Apply a fan-out gate on $n+1$ qubits to Y and $G(1), \ldots, G(n)$, where Y is the control qubit.
3. Apply C_n to $X_1, \ldots, X_n$.
4. Apply a 2-controlled Z gate to $G(j)$, X_j, and W_j for every $1 \leq j \leq n$ in parallel.
5. Apply $C_n^\dagger$ to $X_1, \ldots, X_n$.
6. Apply the gates in Step 2 and Step 1 (in this order).

The fan-out gate can be decomposed into an $O(\log n)$-depth quantum circuit consisting only of CNOT gates. An example of the decomposition is given in Appendix A.2. Moreover, the 2-controlled Z gate can be decomposed into a constant number of the gates in $\mathcal{G}$ [1, 20]. Thus, we can regard E_{2n} as the $O(\log n)$-depth circuit in the case where the elementary gate set is $\mathcal{G}$. The circuit E_{2n} has the desired output probability distribution as follows. The proof can be found in Appendix A.4.

Lemma 3. *For any $x, w \in \{0,1\}^n$, the output of E_{2n} with the input qubits initialized to $|x\rangle|w\rangle$ and output qubit initialized to $|0\rangle$ is 0 with probability $(1 + F_n(x,w))/2$.*

This lemma immediately implies Theorem 2:

Proof of Theorem 2. We modify the above-mentioned probabilistic algorithm for MEE$(p(n), C_n)$ by replacing D_{2n} with E_{2n}. By Lemma 3, the output probability distribution of E_{2n} is the same as that of D_{2n}. Thus, as with the original algorithm, the resulting algorithm solves MEE$(p(n), C_n)$. □

4 Limitations of Uninitialized Ancillary Qubits

4.1 Our Idea for Proving Theorem 3

For any integer $s \geq 1$, an s-controlled Toffoli gate is decomposed into one s-controlled Z gate and two H gates without using any ancillary qubit [9]. Thus, to prove Theorem 3, it suffices to consider an

unbounded Z gate in place of an unbounded Toffoli gate. Our proof of Theorem 3 is an extension of the proof of Bera [3, 2]. We assume on the contrary that there exists a depth-d quantum circuit C_n for PA_n with $p = O(\log n)$ initialized and $q = n^{O(1)}$ uninitialized ancillary qubits such that it consists of one-qubit gates, CNOT gates, unbounded fan-out gates on $(\log n)^{O(1)}$ qubits, and unbounded Z gates, where $d = O((\log n)^{\delta})$ for some constant $0 \leq \delta < 1$. A simple case is when all unbounded Z gates in C_n act on a small number of qubits, such as $O(\log n)$ qubits. In this case, since d is small, a direct application of the previous proof [3] implies that there exists an input qubit of C_n such that the output of C_n does not depend on the input qubit. Thus, C_n cannot compute PA_n since the output of PA_n changes if any one of the n input bits changes. This contradicts the assumption.

The remaining case is when there exists an unbounded Z gate on a large number of qubits. Let $\tilde{C}_n$ be the circuit obtained from C_n by removing all such gates. When C_n does not have any ancillary qubit, Bera [3] showed that C_n is well approximated by $\tilde{C}_n$ in the sense that, when the state of the input qubits is a computational basis state chosen uniformly at random, the output of C_n coincides with that of $\tilde{C}_n$ with high probability. Since C_n computes PA_n, $\tilde{C}_n$ computes PA_n with high probability. Thus, we obtain a contradiction as in the above simple case since all gates in $\tilde{C}_n$ act on a small number of qubits. To apply this idea to our setting where C_n has p initialized and q uninitialized ancillary qubits, we analyze the uninitialized ancillary qubits. Using the assumption that $p = O(\log n)$, we show that C_n with the initialized ancillary qubits (initialized to $|0^p\rangle$) and uninitialized ancillary qubits initialized to $|a\rangle$ for some $a \in \{0,1\}^q$ is well approximated (in the sense described above) by $\tilde{C}_n$ with the same initial state. The former circuit computes PA_n since, by definition, C_n with an arbitrary initial state of the uninitialized ancillary qubits computes PA_n. Thus, the latter circuit computes PA_n with high probability, and we obtain a contradiction as in the above simple case.

4.2 Analysis of a General Circuit and Proof of Theorem 3

We analyze a general depth-d quantum circuit C_n with n input qubits, one output qubit, p initialized ancillary qubits, and q uninitialized ancillary qubits such that it consists of one-qubit gates, CNOT gates, unbounded fan-out gates, and unbounded Z gates. A key property of C_n is described as follows. Bera proved it only when w is constant [3, 2], but the same proof works in the general case.

Lemma 4 ([3, 2])**.** *Let C_n be a depth-d quantum circuit with n input qubits and one output qubit (possibly with ancillary qubits). If all gates in C_n act on at most w qubits, then the output of C_n can depend only on the states of at most w^d input qubits.*

Let $t \geq 2$ be an arbitrary integer and $\mathcal{G}_t$ be the set of all unbounded Z gates in C_n that act on more than or equal to t qubits. We consider only the case where $\mathcal{G}_t \neq \emptyset$ and assume that $\mathcal{G}_t = \{T_1, \ldots, T_k\}$ for some $k \geq 1$, where, for any $1 \leq l \leq k$, if T_l is in layer d of C_n, then T_{l+1} is in layer $d' \geq d$. We decompose C_n into the gates in $\mathcal{G}_t$ and the other parts as depicted in Fig. 2, where C_n computes a Boolean function f_n on n bits and C_n^j is a quantum circuit consisting of gates that are not in $\mathcal{G}_t$ for any $1 \leq j \leq k+1$. Such a decomposition is not unique in general, but the point is to fix a decomposition.

For any $1 \leq l \leq k$, we define a quantum circuit V_l as follows: $V_1 = C_n^1$ and $V_l = C_n^l T_{l-1} V_{l-1}$ for any $2 \leq l \leq k$. Moreover, we define $\Delta_l(x,y,b) = ||T_l V_l |x \circ y \circ b\rangle - V_l |x \circ y \circ b\rangle||$ and $\Delta(x,y,b) = ||C_n|x \circ y \circ b\rangle - \tilde{C}_n|x \circ y \circ b\rangle||$ for any $x \in \{0,1\}^n$, $y \in \{0,1\}$, and $b \in \{0,1\}^{p+q}$. Here, the symbol "$\circ$" represents the concatenation of bit strings, $||\cdot||$ is the Euclidean norm of a vector, i.e., $|||v\rangle|| = \sqrt{\langle v|v\rangle}$ for any vector $|v\rangle$, and $\tilde{C}_n = C_n^{k+1} C_n^k \cdots C_n^2 C_n^1$, i.e., the quantum circuit obtained from C_n by removing $T_1, \ldots, T_k$. Let U_n be a random variable uniformly distributed over $\{0,1\}^n$.

We first evaluate the probability $\Pr[\Delta(U_n,y,b) < \varepsilon]$ using the expected value $\mathrm{E}[\Delta_l(U_n,y,b)^2] = (1/2^n)\sum_{x \in \{0,1\}^n} \Delta_l(x,y,b)^2$ as follows. The proof can be found in Appendix A.5.

Lemma 5. $\Pr[\Delta(U_n,y,b) < \varepsilon] \geq 1 - (k^2/\varepsilon^2)\sum_{l=1}^k \mathrm{E}[\Delta_l(U_n,y,b)^2]$ *for any $\varepsilon > 0$, $y \in \{0,1\}$, and $b \in \{0,1\}^{p+q}$.*

To evaluate the value $\sum_{l=1}^k \mathrm{E}[\Delta_l(U_n,y,b)^2]$, let t_l be the number of qubits on which T_l acts, u_l be the number of the other qubits in C_n, i.e., $u_l = n+p+q+1-t_l$, and $t_{\min} = \min\{t_l | 1 \leq l \leq k\}$. We assume that $V_l|x \circ y \circ b\rangle = \sum_{i \in \{0,1\}^{t_l}} \sum_{j \in \{0,1\}^{u_l}} g_{x \circ y \circ b}^{(l)}(i \circ j)|i \circ j\rangle$ for any $x \in \{0,1\}^n$, $y \in \{0,1\}$, and $b \in \{0,1\}^{p+q}$, where $g_{x \circ y \circ b}^{(l)}(i \circ j) \in \mathbb{C}$. The qubits represented by $i \in \{0,1\}^{t_l}$ correspond to the qubits on which T_l acts. Of course, for any $1 \leq l \leq k$, T_l does not always act on the first t_l qubits in C_n. We

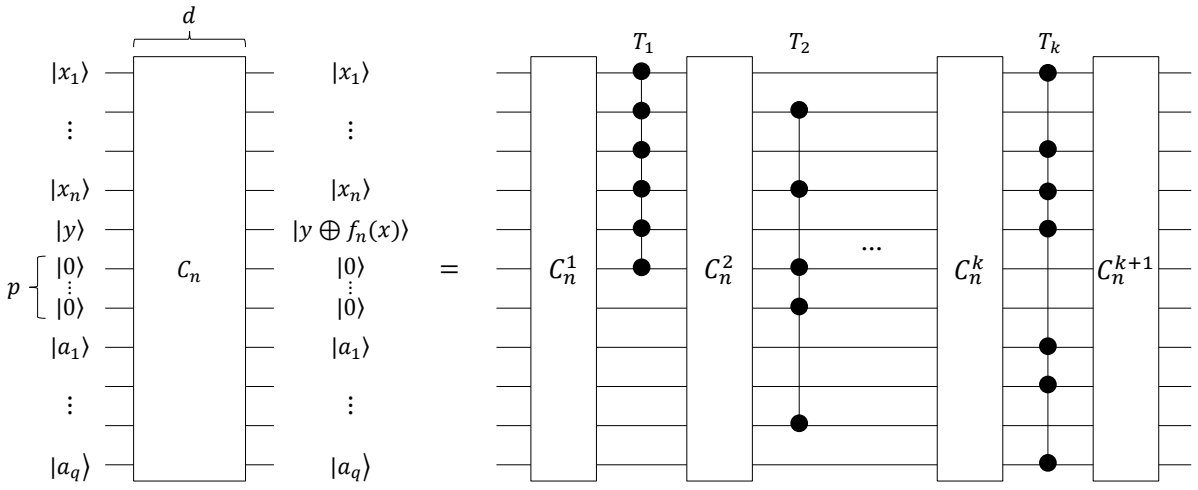

Figure 2: Circuit C_n for f_n and its decomposition, where the initial states of the input qubits, output qubit, and uninitialized ancillary qubits are $|x_1\rangle \cdots |x_n\rangle$, $|y\rangle$, and $|a_1\rangle \cdots |a_q\rangle$, respectively, for any $x = x_1 \cdots x_n \in \{0,1\}^n$, $y \in \{0,1\}$, and $a_1 \cdots a_q \in \{0,1\}^q$. Gates $T_1, \ldots, T_k$ are unbounded Z gates.

therefore need to apply some permutation of all qubits; however, since such a permutation does not affect Lemma 7, which is the key to Theorem 3, we omit it.

We evaluate the above value as follows. The point is that there exists an initial state of the uninitialized ancillary qubits such that this value is small. The proof can be found in Appendix A.6.

Lemma 6. $\sum_{l=1}^{k} \mathrm{E}[\Delta_l(U_n, y, b)^2] \leq k2^{p+q+3}/2^{t_{\min}}$ *for any* $y \in \{0,1\}$ *and* $b \in \{0,1\}^{p+q}$. *Moreover, there exists some* $a \in \{0,1\}^q$ *such that* $\sum_{l=1}^{k} \mathrm{E}[\Delta_l(U_n, 0, 0^p \circ a)^2] \leq k2^{p+3}/2^{t_{\min}}$.

Lemmas 5 and 6 immediately imply the following evaluation:

Lemma 7. $\Pr[\Delta(U_n, y, b) < \varepsilon] \geq 1 - k^3 2^{p+q+3}/(\varepsilon^2 2^{t_{\min}})$ *for any* $\varepsilon > 0$, $y \in \{0,1\}$, *and* $b \in \{0,1\}^{p+q}$. *In particular, there exists some* $a \in \{0,1\}^q$ *such that* $\Pr[\Delta(U_n, 0, 0^p \circ a) < \varepsilon] \geq 1 - k^3 2^{p+3}/(\varepsilon^2 2^{t_{\min}})$.

Lemmas 4 and 7 imply Theorem 3 as follows. We assume on the contrary that there exists a depth-d quantum circuit C_n for PA_n with $p = O(\log n)$ initialized and $q = n^{O(1)}$ uninitialized ancillary qubits such that it consists of one-qubit gates, CNOT gates, unbounded fan-out gates on $(\log n)^{O(1)}$ qubits, and unbounded Z gates, where $d = O((\log n)^{\delta})$ for some constant $0 \leq \delta < 1$. Let t be a sufficiently large number satisfying that $t = \Theta(\log(n + p + q + 1))$. When $\mathcal{G}_t = \emptyset$, all gates in C_n act on at most $w = (\log n)^{O(1)}$ qubits. By Lemma 4, the output of C_n can depend only on the states of at most $w^d = o(n)$ input qubits. As described in Section 4.1, C_n cannot compute PA_n, which is a contradiction. When $\mathcal{G}_t \neq \emptyset$, we apply the above analysis of a general circuit. By Lemma 7 with $\varepsilon = 0.1$, there exists some $a \in \{0,1\}^q$ such that $\Pr[\Delta(U_n, 0, 0^p \circ a) < 0.1]$ is polynomially (in n) close to 1. This means that, when n is sufficiently large, there are more than 2^{n-1} elements $x \in \{0,1\}^n$ such that $\Delta(x, 0, 0^p \circ a) < 0.1$. Since C_n computes PA_n, it is easy to show that, for any such x, the output of $\tilde{C}_n|x \circ 0 \circ 0^p \circ a\rangle$ is $\mathrm{PA}_n(x)$ with probability of at least $1 - 0.1^2 = 0.99$. On the other hand, Lemma 4 implies that, for at most 2^{n-1} elements $x \in \{0,1\}^n$, the output of $\tilde{C}_n|x \circ 0 \circ 0^p \circ a\rangle$ is $\mathrm{PA}_n(x)$ with probability greater than 0.5. This is a contradiction. The details can be found in Appendix A.7.

5 Open Problems

Interesting challenges would be to further study the computational power of shallow quantum circuits with uninitialized ancillary qubits. We give some examples of such problems:

- Can we decrease the depth or the number of qubits on which fan-out gates act in Theorem 1?
- What functions (not dealt with in this paper) can or cannot be computed by shallow quantum circuits with uninitialized ancillary qubits?
- What is the relationship between the computational power of general classical/quantum circuits and that of shallow quantum circuits with uninitialized ancillary qubits?

References

[1] A. Barenco, C. H. Bennett, R. Cleve, D. P. DiVincenzo, N. Margolus, P. Shor, T. Sleator, J. A. Smolin, and H. Weinfurter. Elementary gates for quantum computation. *Physical Review A*, 52(5):3457–3467, 1995.

[2] D. Bera. *Quantum circuits: power and limitations*. PhD thesis, Boston University, 2010.

[3] D. Bera. A lower bound method for quantum circuits. *Information Processing Letters*, 111(15):723–726, 2011.

[4] H. Buhrman, R. Cleve, M. Koucký, B. Loff, and F. Speelman. Computing with a full memory: catalytic space. In *Proceedings of the 46th ACM Symposium on Theory of Computing (STOC)*, pages 857–866, 2014.

[5] R. Cleve and J. Watrous. Fast parallel circuits for the quantum Fourier transform. In *Proceedings of the 41st IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 526–536, 2000.

[6] D. Coppersmith. An approximate Fourier transform useful in quantum factoring, 2002. arXiv:quant-ph/0201067.

[7] A. Cosentino, R. Kothari, and A. Paetznick. Dequantizing read-once quantum formulas. In *Proceedings of the 8th Conference on the Theory of Quantum Computation, Communication and Cryptography (TQC)*, volume 22 of *Leibniz International Proceedings in Informatics*, pages 80–92, 2013.

[8] D. P. DiVincenzo. The physical implementation of quantum computation. *Fortschritte der Physik*, 48(9–11):771–783, 2000.

[9] M. Fang, S. Fenner, F. Green, S. Homer, and Y. Zhang. Quantum lower bounds for fanout. *Quantum Information and Computation*, 6(1):46–57, 2006.

[10] S. Fenner, F. Green, S. Homer, and Y. Zhang. Bounds on the power of constant-depth quantum circuits. In *Proceedings of Fundamentals of Computation Theory (FCT)*, volume 3623 of *Lecture Notes in Computer Science*, pages 44–55, 2005.

[11] O. Goldreich. *Foundations of Cryptography: Volume I, Basic Tools*. Cambridge University Press, 2001.

[12] F. Green, S. Homer, C. Moore, and C. Pollett. Counting, fanout, and the complexity of quantum ACC. *Quantum Information and Computation*, 2(1):35–65, 2002.

[13] P. Høyer and R. Špalek. Quantum fan-out is powerful. *Theory of Computing*, 1(5):81–103, 2005.

[14] A. Jakoby, M. Liśkiewicz, and R. Reischuk. Private computations in networks: topology versus randomness. In *Proceedings of the 20th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 2607 of *Lecture Notes in Computer Science*, pages 121–132, 2003.

[15] S. Jukna. *Boolean Function Complexity: Advances and Frontiers*. Springer, 2012.

[16] E. Knill and R. Laflamme. Power of one bit of quantum information. *Physical Review Letters*, 81(25):5672–5675, 1998.

[17] G. De las Cuevas, W. Dür, M. van den Nest, and M. A. Martin-Delgado. Quantum algorithms for classical lattice models. *New Journal of Physics*, 13(093021), 2011.

[18] C. Moore and M. Nilsson. Parallel quantum computation and quantum codes. *SIAM Journal on Computing*, 31(3):799–815, 2001.

[19] X. Ni and M. van den Nest. Commuting quantum circuits: efficient classical simulations versus hardness results. *Quantum Information and Computation*, 13(1&2):54–72, 2013.

[20] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

[21] M. Saeedi and M. Pedram. Linear-depth quantum circuits for n-qubit Toffoli gates with no ancilla. *Physical Review A*, 87(062318), 2013.

[22] Y. Takahashi and N. Kunihiro. A quantum circuit for Shor's factoring algorithm using $2n+2$ qubits. *Quantum Information and Computation*, 6(2):184–192, 2006.

[23] Y. Takahashi and S. Tani. Collapse of the hierarchy of constant-depth exact quantum circuits. *Computational Complexity*, 25(4):849–881, 2016.

[24] Y. Takahashi, T. Yamazaki, and K. Tanaka. Hardness of classically simulating quantum circuits with unbounded Toffoli and fan-out gates. *Quantum Information and Computation*, 14(13&14):1149–1164, 2014.

[25] H. Vollmer. *Introduction to Circuit Complexity*. Springer, 1999.

A Proofs

A.1 Proof of Lemma 1

Proof. As an example, Stages 1 and 2 with $n = 3$ are depicted in Figs. 3 and 4, respectively. The states of $X_1, \ldots, X_n$ stay unchanged during the computation since the qubits are used only for control qubits. We fix an arbitrary $1 \leq k \leq m$ and show the lemma by induction on s. We first consider the base case, $s = 1$. Steps 1–2 transform the initial state of $I(k)$ and $B_1(k,1), \ldots, B_n(k,1)$ as follows:

$$\begin{aligned}|0\rangle|b_1(k,1)\rangle\cdots|b_n(k,1)\rangle &\mapsto \frac{1}{\sqrt{2}}|0\rangle|b_1(k,1)\rangle\cdots|b_n(k,1)\rangle + \frac{1}{\sqrt{2}}|1\rangle|b_1(k,1)\rangle\cdots|b_n(k,1)\rangle \\ &\mapsto \frac{1}{\sqrt{2}}|0\rangle|b_1(k,1)\rangle\cdots|b_n(k,1)\rangle + \frac{1}{\sqrt{2}}|1\rangle|b_1(k,1)\oplus 1\rangle\cdots|b_n(k,1)\oplus 1\rangle,\end{aligned}$$

where the symbol "$\mapsto$" corresponds to one step. Step 3 does nothing and Step 4 transforms this state into the state

$$\frac{1}{\sqrt{2}}|0\rangle|b_1(k,1)\rangle\cdots|b_n(k,1)\rangle + \frac{e^{\frac{2\pi i}{2^k}\alpha(k,1)}}{\sqrt{2}}|1\rangle|b_1(k,1)\oplus 1\rangle\cdots|b_n(k,1)\oplus 1\rangle, \tag{1}$$

where we ignore the global phase and

$$\alpha(k,1) = -\sum_{j=1}^{n}((b_j(k,1)\oplus 1) - b_j(k,1))a_j(k) = -\sum_{j=1}^{n}(1-2b_j(k,1))a_j(k).$$

Step 5 transforms the states of $A_1(k), \ldots, A_n(k)$ into the states $|x_1 \oplus a_1(k)\rangle, \ldots, |x_n \oplus a_n(k)\rangle$, respectively. Step 6 transforms state (1) into the state

$$\frac{1}{\sqrt{2}}|0\rangle|b_1(k,1)\rangle\cdots|b_n(k,1)\rangle + \frac{e^{\frac{2\pi i}{2^k}(\alpha(k,1)+\beta(k,1))}}{\sqrt{2}}|1\rangle|b_1(k,1)\oplus 1\rangle\cdots|b_n(k,1)\oplus 1\rangle,$$

where

$$\beta(k,1) = \sum_{j=1}^{n}((b_j(k,1)\oplus 1) - b_j(k,1))(x_j \oplus a_j(k)) = \sum_{j=1}^{n}(1-2b_j(k,1))(x_j \oplus a_j(k)).$$

Thus,

$$\begin{aligned}\alpha(k,1)+\beta(k,1) &= \sum_{j=1}^{n}(1-2b_j(k,1))((x_j \oplus a_j(k)) - a_j(k)) = \sum_{j=1}^{n}(-1)^{b_j(k,1)}(-1)^{a_j(k)}x_j \\ &= \sum_{j=1}^{n}x_j(1-2(a_j(k)\oplus b_j(k,1))) = |x| - 2\sum_{j=1}^{n}x_j(a_j(k)\oplus b_j(k,1)) = \gamma(k,1).\end{aligned}$$

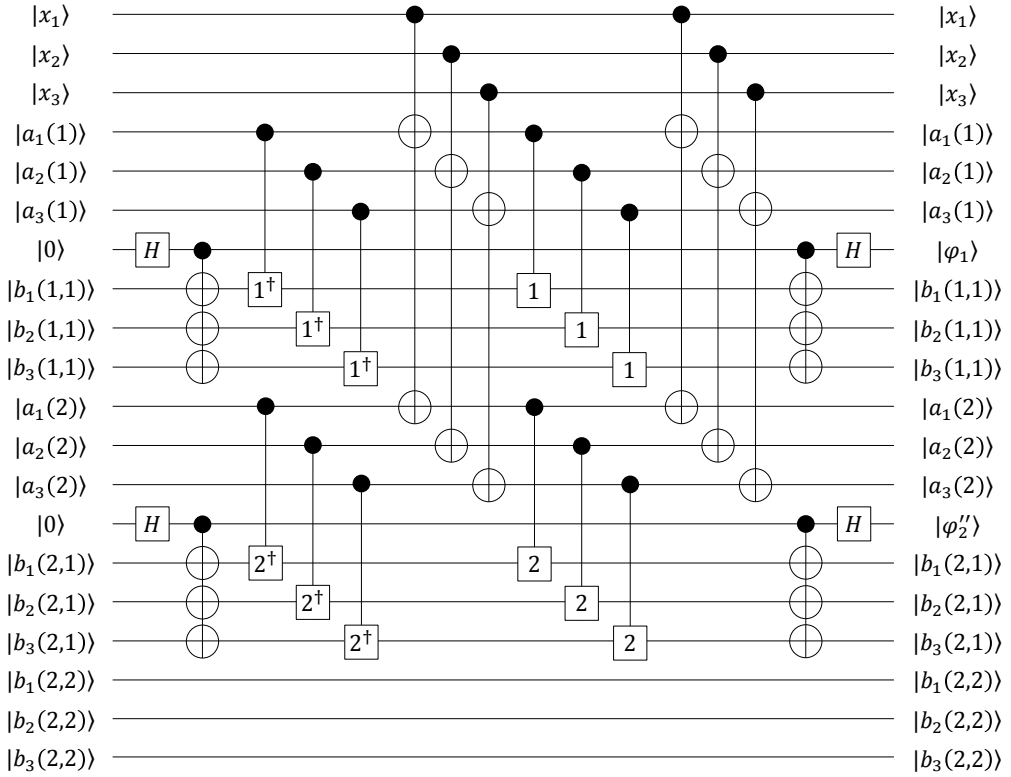

Figure 3: Stage 1. For any integer $t \geq 1$, the gate represented as $t^\dagger$ is the inverse of the $Z(2\pi/2^t)$ gate, i.e., $Z(-2\pi/2^t)$ gate.

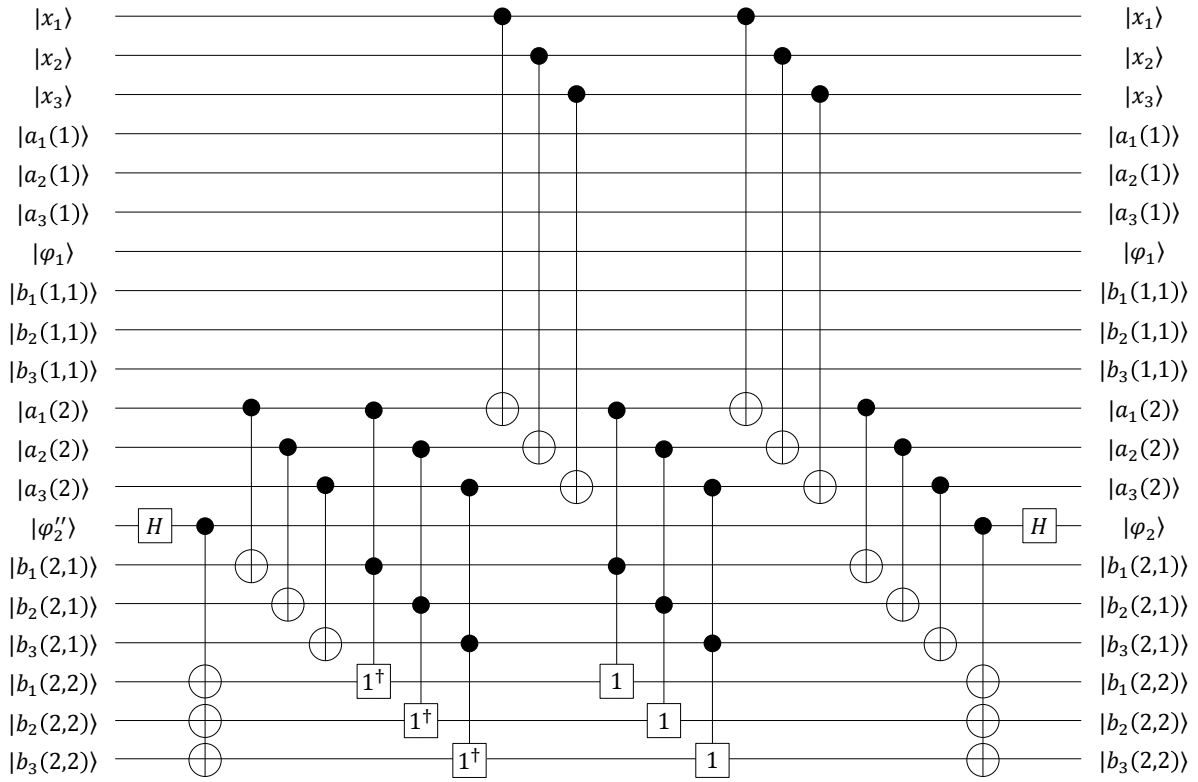

Figure 4: Stage 2.

Steps 7–9 transform the state of all the qubits other than $I(k)$ into their initial state. The state of $I(k)$ after Step 10 is

$$\frac{|+\rangle + e^{\frac{2\pi i}{2^k}\gamma(k,1)}|-\rangle}{\sqrt{2}}. \tag{2}$$

The state $|\varphi_2''\rangle$ in Figs. 3 and 4 is state (2) when $n = 3$ and $k = 2$. This completes the proof of the base case. In particular, the above proof implies that Lemma 1 holds when $k = 1$, and thus we assume that $k \geq 2$ in the following.

We fix an arbitrary $1 \leq s \leq k-1$ and assume that the lemma holds for s. Thus, after Stage s, the state of $I(k)$ is

$$\frac{|+\rangle + e^{\frac{2\pi i}{2^k}\gamma(k,s)}|-\rangle}{\sqrt{2}}.$$

Moreover, for any qubit other than the initialized ancillary qubits, the state of the qubit is the same as its initial state. We apply Stage $s+1 \leq k$. Step 1 transforms the initial state of $I(k)$ and $B_1(k,s+1),\ldots,B_n(k,s+1)$ as follows:

$$\begin{aligned}
&\frac{|+\rangle + e^{\frac{2\pi i}{2^k}\gamma(k,s)}|-\rangle}{\sqrt{2}}|b_1(k,s+1)\rangle\cdots|b_n(k,s+1)\rangle \\
&\mapsto \frac{1}{\sqrt{2}}|0\rangle|b_1(k,s+1)\rangle\cdots|b_n(k,s+1)\rangle + \frac{e^{\frac{2\pi i}{2^k}\gamma(k,s)}}{\sqrt{2}}|1\rangle|b_1(k,s+1)\rangle\cdots|b_n(k,s+1)\rangle.
\end{aligned}$$

Step 2 transforms this state into the state

$$\frac{1}{\sqrt{2}}|0\rangle|b_1(k,s+1)\rangle\cdots|b_n(k,s+1)\rangle + \frac{e^{\frac{2\pi i}{2^k}\gamma(k,s)}}{\sqrt{2}}|1\rangle|b_1(k,s+1)\oplus 1\rangle\cdots|b_n(k,s+1)\oplus 1\rangle. \tag{3}$$

Step 3 transforms the states of $B_j(k,1),\ldots,B_j(k,s)$ into the states

$$|a_j(k)\oplus b_j(k,1)\rangle,\ldots,|a_j(k)\oplus b_j(k,s)\rangle,$$

respectively, for any $1 \leq j \leq n$. Step 4 transforms state (3) into the state

$$\frac{1}{\sqrt{2}}|0\rangle|b_1(k,s+1)\rangle\cdots|b_n(k,s+1)\rangle + \frac{e^{\frac{2\pi i}{2^k}(\gamma(k,s)+\alpha(k,s+1))}}{\sqrt{2}}|1\rangle|b_1(k,s+1)\oplus 1\rangle\cdots|b_n(k,s+1)\oplus 1\rangle, \tag{4}$$

where

$$\begin{aligned}
\alpha(k,s+1) &= -2^s\sum_{j=1}^{n}((b_j(k,s+1)\oplus 1) - b_j(k,s+1))a_j(k)\bigwedge_{l=1}^{s}(a_j(k)\oplus b_j(k,l)) \\
&= -2^s\sum_{j=1}^{n}(1-2b_j(k,s+1))a_j(k)\bigwedge_{l=1}^{s}(a_j(k)\oplus b_j(k,l)).
\end{aligned}$$

Step 5 transforms the states of $A_1(k),\ldots,A_n(k)$ into the states $|x_1\oplus a_1(k)\rangle,\ldots,|x_n\oplus a_n(k)\rangle$, respectively. Step 6 transforms state (4) into the state

$$\frac{1}{\sqrt{2}}|0\rangle|b_1(k,s+1)\rangle\cdots|b_n(k,s+1)\rangle + \frac{e^{\frac{2\pi i}{2^k}(\gamma(k,s)+\alpha(k,s+1)+\beta(k,s+1))}}{\sqrt{2}}|1\rangle|b_1(k,s+1)\oplus 1\rangle\cdots|b_n(k,s+1)\oplus 1\rangle,$$

where

$$\begin{aligned}
\beta(k,s+1) &= 2^s\sum_{j=1}^{n}((b_j(k,s+1)\oplus 1) - b_j(k,s+1))(x_j\oplus a_j(k))\bigwedge_{l=1}^{s}(a_j(k)\oplus b_j(k,l)) \\
&= 2^s\sum_{j=1}^{n}(1-2b_j(k,s+1))(x_j\oplus a_j(k))\bigwedge_{l=1}^{s}(a_j(k)\oplus b_j(k,l)).
\end{aligned}$$

This implies that

$$\begin{aligned}
\alpha(k,s+1)+\beta(k,s+1) &= 2^s\sum_{j=1}^{n}(1-2b_j(k,s+1))((x_j\oplus a_j(k))-a_j(k))\bigwedge_{l=1}^{s}(a_j(k)\oplus b_j(k,l))\\
&= 2^s\sum_{j=1}^{n}(-1)^{b_j(k,s+1)}(-1)^{a_j(k)}x_j\bigwedge_{l=1}^{s}(a_j(k)\oplus b_j(k,l))\\
&= 2^s\sum_{j=1}^{n}x_j(1-2(a_j(k)\oplus b_j(k,s+1)))\bigwedge_{l=1}^{s}(a_j(k)\oplus b_j(k,l))\\
&= 2^s\sum_{j=1}^{n}x_j\bigwedge_{l=1}^{s}(a_j(k)\oplus b_j(k,l))-2^{s+1}\sum_{j=1}^{n}x_j\bigwedge_{l=1}^{s+1}(a_j(k)\oplus b_j(k,l))\\
&= |x|-\gamma(k,s)-2^{s+1}\sum_{j=1}^{n}x_j\bigwedge_{l=1}^{s+1}(a_j(k)\oplus b_j(k,l))\\
&= \gamma(k,s+1)-\gamma(k,s).
\end{aligned}$$

Thus, $\gamma(k,s)+\alpha(k,s+1)+\beta(k,s+1)=\gamma(k,s+1)$. Steps 7–9 transform the state of all the qubits other than $I(k)$ into their initial state. The state of $I(k)$ after Step 10 is

$$\frac{|+\rangle+e^{\frac{2\pi i}{2^k}\gamma(k,s+1)}|-\rangle}{\sqrt{2}}.$$

Therefore, the lemma holds for $s+1$ as desired. We note that

$$e^{\frac{2\pi i}{2^k}\gamma(k,k)}=e^{\frac{2\pi i}{2^k}(|x|-2^k\sum_{j=1}^{n}x_j\bigwedge_{l=1}^{k}(a_j(k)\oplus b_j(k,l)))}=e^{\frac{2\pi i}{2^k}|x|}.$$

This completes the proof of the lemma. □

A.2 Proof of Lemma 2

Proof. We first consider the case where the elementary gate set is $\mathcal{G}$ with an unbounded Toffoli gate on $O(\log n)$ qubits and an unbounded fan-out gate. For any $k\geq 2$, a k-controlled $Z'(2\pi/2^t)$ gate is decomposed into one 1-controlled $Z'(2\pi/2^t)$ gate, two 1-controlled $Z'(-2\pi/2^{t+1})$ gates, and two $(k-1)$-controlled Toffoli gates without using any new ancillary qubit [1]. For example, an 8-controlled $Z'(2\pi/2^t)$ gate is decomposed as depicted in Fig. 5(a). A k-controlled $Z'(-2\pi/2^t)$ gate is decomposed similarly. Moreover, a 1-controlled $Z'(\pm 2\pi/2^r)$ gate for any integer $r\geq 1$ is decomposed into two one-qubit gates of the form $Z'(\pm 2\pi/2^{r'})$ for some integer $r'\geq 1$ and two CNOT gates [1]. By this decomposition, we can regard Q_n as the circuit in the case where the elementary gate set is $\mathcal{G}$ with an unbounded Toffoli gate on $O(\log n)$ qubits and an unbounded fan-out gate. The depth of each stage is constant and thus the depth of the whole circuit, which consists of $m=O(\log n)$ stages, is $O(\log n)$. It is obvious that the circuit uses $m=O(\log n)$ initialized ancillary qubits and $nm(m+3)/2=O(n(\log n)^2)$ uninitialized ancillary qubits.

A k-controlled Toffoli gate is decomposed into an $O(k)$-depth quantum circuit with an uninitialized ancillary qubit such that it consists only of H gates, $Z(\pm\pi/4)$ gates, and CNOT gates [1, 20]. Moreover, by using the equivalence of a parity gate and a fan-out gate [12], a fan-out gate on $k+1$ qubits is decomposed into an $O(\log k)$-depth quantum circuit without using any new ancillary qubit such that it consists only of CNOT gates. An example of such a circuit with $k=8$ is depicted in Fig. 5(b). By this decomposition, we can regard Q_n as the circuit in the case where the elementary gate set is $\mathcal{G}$. The depth of each stage is $O(\log n)$ since an unbounded fan-out gate acts on at most $n+1=O(n)$ qubits and an unbounded Toffoli gate acts on at most $m+1=O(\log n)$ qubits. Thus, the depth of the whole circuit is $O(m\log n)=O((\log n)^2)$. Moreover, as in the above case, the circuit uses m initialized ancillary qubits and $nm(m+3)/2$ uninitialized ancillary qubits. □

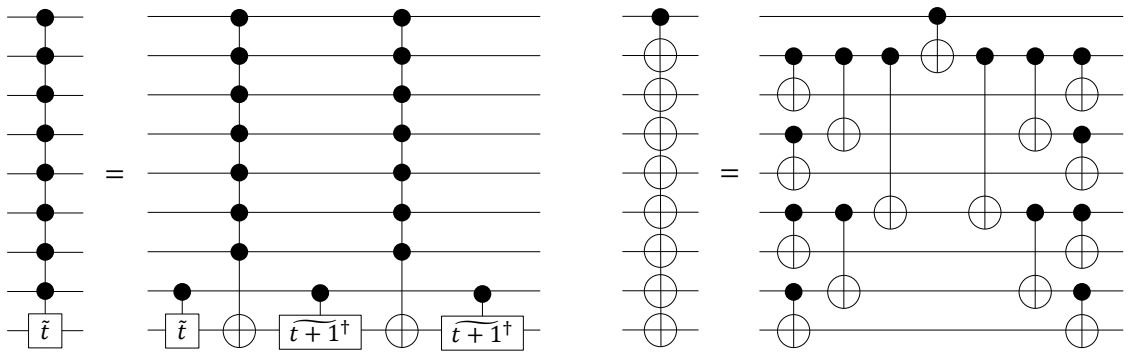

(a) (b)

Figure 5: (a): Decomposition of an 8-controlled $Z'(2\pi/2^t)$ gate, where the gates represented as $\tilde{t}$ and $\tilde{t}^{\dagger}$ are a $Z'(2\pi/2^t)$ gate and its inverse, i.e., a $Z'(-2\pi/2^t)$ gate, respectively, for any integer $t \geq 1$. (b): Decomposition of a fan-out gate on nine qubits.

A.3 Proof of Theorem 1

Proof. Let $f_n \in \mathcal{E}$ and $x = x_1 \cdots x_n \in \{0,1\}^n$ be an input. By the definition of $\mathcal{E}$, there exists an $O(\log n)$-size (and thus $O(\log n)$-depth) quantum circuit R_m with $m = \lceil \log(n+1) \rceil$ input qubits, one output qubit, and $r = O(\log n)$ initialized ancillary qubits such that it consists of the gates in $\mathcal{G}$, and it computes $f_n(x)$ when we are given the binary representation of $|x|$.

We prepare n input qubits $X_1, \ldots, X_n$, one output qubit Y, $p = m + r = O(\log n)$ initialized ancillary qubits $I(1), \ldots, I(p)$, and $q = nm(m+3)/2 = O(n(\log n)^2)$ uninitialized ancillary qubits $C(1), \ldots, C(q)$. The initial states of $X_1, \ldots, X_n$ are $|x_1\rangle, \cdots, |x_n\rangle$, respectively, and the initial state of Y is $|y\rangle$ for any $y \in \{0,1\}$. Each of $I(1), \ldots, I(p)$ is initialized to $|0\rangle$ and the initial state of the uninitialized ancillary qubits is arbitrary. We consider the following circuit:

1. Apply Q_n to $X_1, \ldots, X_n$, $I(1), \ldots, I(m)$, and $C(1), \ldots, C(q)$, where $X_1, \ldots, X_n$, $I(1), \ldots, I(m)$, and $C(1), \ldots, C(q)$ are used as the n input qubits, m initialized ancillary qubits, and q uninitialized ancillary qubits of Q_n, respectively.
2. Apply a H gate to $I(k)$ for every $1 \leq k \leq m$ in parallel.
3. Apply a quantum circuit for $\mathrm{QFT}_{2^m}^{\dagger}$ to $I(1), \ldots, I(m)$.
4. Apply R_m to $I(1), \ldots, I(p)$ and Y, where $I(1), \ldots, I(m)$, Y, and $I(m+1), \ldots, I(p)$ are used as the m input qubits, one output qubit, and r initialized ancillary qubits of R_m, respectively.
5. Apply the inverse of the circuit in Step 3.
6. Same as Step 2.
7. Apply the inverse of the circuit in Step 1.

By Lemma 1, Step 1 transforms the initial state of $I(1), \ldots, I(m)$ into the state $\bigotimes_{k=1}^{m} |\varphi_k\rangle$. After Step 1, the state of the other qubits is the same as their initial state. Step 2 transforms this state into the state $\bigotimes_{k=1}^{m} H|\varphi_k\rangle$. As described in Section 2.2, Step 3 transforms this state into the state $\bigotimes_{k=1}^{m} |s_k\rangle$, which is the binary representation of $|x|$. Step 4 transforms the initial state of Y into the state $|y \oplus f_n(x)\rangle$. After Step 4, the state of the other qubits is the same as their state after Step 3. Steps 5–7 return the state of $I(1), \ldots, I(m)$ into their initial state. Therefore, the above circuit computes f_n.

It is known that there exists an $O(\log n)$-depth $O((\log n)^2)$-size quantum circuit for QFT_{2^m} with no ancillary qubits such that it consists of H gates and 1-controlled $Z(2\pi/2^t)$ gates for any integer $t \geq 1$ [6, 20]. In the above context, a direct calculation immediately shows that we can replace a 1-controlled $Z(2\pi/2^t)$ gate with a 1-controlled $Z'(2\pi/2^t)$ gate, which is decomposed into two one-qubit gates of the form $Z'(\pm 2\pi/2^r)$ and two CNOT gates [1]. Thus, when the elementary gate set is $\mathcal{G}$, the depth of the circuit consisting only of Steps 2–6 is $O(\log n)$. In this case, by Lemma 2, the depth of Q_n

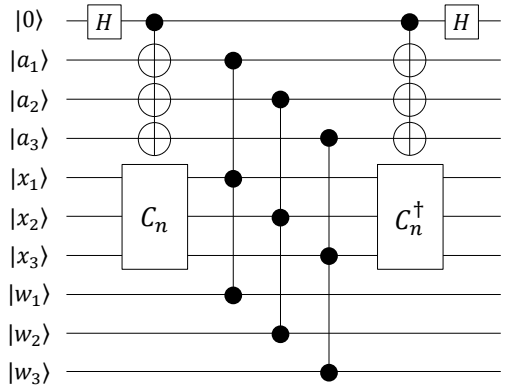

Figure 6: Circuit E_{2n} with $n = 3$. The top qubit is the output qubit.

is $O((\log n)^2)$ and thus the depth of the whole circuit is $O((\log n)^2)$ as desired. When the elementary gate set further includes an unbounded Toffoli gate on $O(\log n)$ qubits and an unbounded fan-out gate, by Lemma 2, the depth of Q_n is $O(\log n)$ and thus the depth of the whole circuit is $O(\log n)$ as desired. This completes the proof of Theorem 1. □

A.4 Proof of Lemma 3

Proof. As an example, E_{2n} with $n = 3$ is depicted in Fig. 6. We assume that the initial states of $G(1), \ldots, G(n)$ are $|a_1\rangle, \ldots, |a_n\rangle$, respectively, for any $a = a_1 \cdots a_n \in \{0,1\}^n$. The states of $W_1, \ldots, W_n$ stay unchanged during the computation since the qubits are used only for control qubits. Thus, we consider the state of the remaining qubits $Y, G(1), \ldots, G(n), X_1, \ldots, X_n$. Steps 1–2 transform their initial state $|0\rangle|a\rangle|x\rangle$ as follows:

$$|0\rangle|a\rangle|x\rangle \mapsto \frac{1}{\sqrt{2}}|0\rangle|a\rangle|x\rangle + \frac{1}{\sqrt{2}}|1\rangle|a\rangle|x\rangle \mapsto \frac{1}{\sqrt{2}}|0\rangle|a_1\rangle\cdots|a_n\rangle|x\rangle + \frac{1}{\sqrt{2}}|1\rangle|a_1 \oplus 1\rangle\cdots|a_n \oplus 1\rangle|x\rangle,$$

where the symbol "$\mapsto$" corresponds to one step. Step 3 transforms this state into the state

$$\frac{1}{\sqrt{2}}|0\rangle|a_1\rangle\cdots|a_n\rangle(C_n|x\rangle) + \frac{1}{\sqrt{2}}|1\rangle|a_1 \oplus 1\rangle\cdots|a_n \oplus 1\rangle(C_n|x\rangle).$$

Step 4 transforms this state into the state

$$\frac{1}{\sqrt{2}}|0\rangle|a_1\rangle\cdots|a_n\rangle\left\{\left(\bigotimes_{j=1}^{n} Z_j^{a_j w_j}\right) C_n|x\rangle\right\} + \frac{1}{\sqrt{2}}|1\rangle|a_1 \oplus 1\rangle\cdots|a_n \oplus 1\rangle\left\{\left(\bigotimes_{j=1}^{n} Z_j^{(a_j \oplus 1) w_j}\right) C_n|x\rangle\right\}.$$

Step 5 transforms this state into the state

$$\frac{1}{\sqrt{2}}|0\rangle|a_1\rangle\cdots|a_n\rangle\left\{C_n^\dagger\left(\bigotimes_{j=1}^{n} Z_j^{a_j w_j}\right) C_n|x\rangle\right\}$$
$$+ \frac{1}{\sqrt{2}}|1\rangle|a_1 \oplus 1\rangle\cdots|a_n \oplus 1\rangle\left\{C_n^\dagger\left(\bigotimes_{j=1}^{n} Z_j^{(a_j \oplus 1) w_j}\right) C_n|x\rangle\right\}.$$

Step 6 transforms this state into the state

$$\frac{|0\rangle + |1\rangle}{2}|a\rangle\left\{C_n^\dagger\left(\bigotimes_{j=1}^{n} Z_j^{a_j w_j}\right) C_n|x\rangle\right\} + \frac{|0\rangle - |1\rangle}{2}|a\rangle\left\{C_n^\dagger\left(\bigotimes_{j=1}^{n} Z_j^{(a_j \oplus 1) w_j}\right) C_n|x\rangle\right\},$$

which is equal to

$$\frac{1}{2}|0\rangle|a\rangle\left\{C_n^\dagger\left(\bigotimes_{j=1}^{n} Z_j^{a_j w_j} + \bigotimes_{j=1}^{n} Z_j^{(a_j \oplus 1) w_j}\right) C_n|x\rangle\right\}$$
$$+ \frac{1}{2}|1\rangle|a\rangle\left\{C_n^\dagger\left(\bigotimes_{j=1}^{n} Z_j^{a_j w_j} - \bigotimes_{j=1}^{n} Z_j^{(a_j \oplus 1) w_j}\right) C_n|x\rangle\right\}.$$

Thus, using the relationships $Z = Z^\dagger$ and $Z^{a_j w_j + (a_j \oplus 1) w_j} = Z^{w_j}$ for any $a_j, w_j \in \{0,1\}$, we can compute the probability that the output of E_{2n} is 0, denoted by $p(0)$, as follows:

$$p(0) = \frac{1}{4}\langle x|C_n^\dagger \left(\bigotimes_{j=1}^n Z_j^{a_j w_j} + \bigotimes_{j=1}^n Z_j^{(a_j \oplus 1) w_j} \right)^2 C_n|x\rangle = \frac{1}{2}\langle x|C_n^\dagger \left(I + \bigotimes_{j=1}^n Z_j^{a_j w_j + (a_j \oplus 1) w_j} \right) C_n|x\rangle$$
$$= \frac{1}{2}\langle x|C_n^\dagger \left(I + \bigotimes_{j=1}^n Z_j^{w_j} \right) C_n|x\rangle = \frac{1 + \langle x|C_n^\dagger Z(w) C_n|x\rangle}{2} = \frac{1 + F_n(x,w)}{2}.$$

□

A.5 Proof of Lemma 5

Proof. Fix an arbitrary $\varepsilon > 0$ and let $\varepsilon_1 = \varepsilon/k$. For any $x \in \{0,1\}^n$, $y \in \{0,1\}$, and $b \in \{0,1\}^{p+q}$, by the triangle inequality [20] and the definition of V_l, it holds that

$$\begin{aligned}
\Delta(x,y,b) &= ||C_n^{k+1} T_k V_k |x \circ y \circ b\rangle - C_n^{k+1} C_n^k \cdots C_n^2 C_n^1 |x \circ y \circ b\rangle|| \\
&\leq ||C_n^{k+1} T_k V_k |x \circ y \circ b\rangle - C_n^{k+1} V_k |x \circ y \circ b\rangle|| \\
&\quad + ||C_n^{k+1} C_n^k T_{k-1} V_{k-1} |x \circ y \circ b\rangle - C_n^{k+1} C_n^k V_{k-1} |x \circ y \circ b\rangle|| \\
&\quad + ||C_n^{k+1} C_n^k C_n^{k-1} T_{k-2} V_{k-2} |x \circ y \circ b\rangle - C_n^{k+1} C_n^k C_n^{k-1} V_{k-2} |x \circ y \circ b\rangle|| \\
&\quad \cdots \\
&\quad + ||C_n^{k+1} C_n^k C_n^{k-1} \cdots C_n^2 T_1 V_1 |x \circ y \circ b\rangle - C_n^{k+1} C_n^k C_n^{k-1} \cdots C_n^2 V_1 |x \circ y \circ b\rangle|| \\
&= \sum_{l=1}^k \Delta_l(x,y,b).
\end{aligned}$$

This implies that, for any $y \in \{0,1\}$ and $b \in \{0,1\}^{p+q}$,

$$\Pr[\Delta(U_n,y,b) < k\varepsilon_1] \geq \Pr[\bigwedge_{l=1}^k \Delta_l(U_n,y,b) < \varepsilon_1]. \tag{5}$$

By the union bound and the Markov's inequality [11],

$$\Pr[\bigvee_{l=1}^k \Delta_l(U_n,y,b) \geq \varepsilon_1] \leq \sum_{l=1}^k \Pr[\Delta_l(U_n,y,b) \geq \varepsilon_1] = \sum_{l=1}^k \Pr[\Delta_l(U_n,y,b)^2 \geq \varepsilon_1^2]$$
$$\leq \frac{1}{\varepsilon_1^2} \sum_{l=1}^k \mathrm{E}[\Delta_l(U_n,y,b)^2],$$

which implies that

$$\Pr[\bigwedge_{l=1}^k \Delta_l(U_n,y,b) < \varepsilon_1] \geq 1 - \frac{1}{\varepsilon_1^2} \sum_{l=1}^k \mathrm{E}[\Delta_l(U_n,y,b)^2].$$

This relationship with (5) implies that

$$\Pr[\Delta(U_n,y,b) < k\varepsilon_1] \geq 1 - \frac{1}{\varepsilon_1^2} \sum_{l=1}^k \mathrm{E}[\Delta_l(U_n,y,b)^2],$$

which immediately implies the desired relationship since $\varepsilon_1 = \varepsilon/k$. □

A.6 Proof of Lemma 6

Proof. First, we show that, for any $1 \leq l \leq k$, $y \in \{0,1\}$, and $b \in \{0,1\}^{p+q}$,

$$\mathrm{E}[\Delta_l(U_n,y,b)^2] = \frac{4}{2^n} \sum_{j \in \{0,1\}^{u_l}} \sum_{x \in \{0,1\}^n} |g_{x \circ y \circ b}^{(l)}(1^{t_l} \circ j)|^2.$$

This can be shown by the following simple calculation. For any $1 \leq l \leq k$, $x \in \{0,1\}^n$, $y \in \{0,1\}$, and $b \in \{0,1\}^{p+q}$,

$$T_l V_l |x \circ y \circ b\rangle - V_l |x \circ y \circ b\rangle = -2 \sum_{j \in \{0,1\}^{u_l}} g^{(l)}_{x \circ y \circ b}(1^{t_l} \circ j) |1^{t_l} \circ j\rangle.$$

This implies that

$$\Delta_l(x,y,b)^2 = 4 \sum_{j \in \{0,1\}^{u_l}} |g^{(l)}_{x \circ y \circ b}(1^{t_l} \circ j)|^2.$$

Thus,

$$\mathrm{E}[\Delta_l(U_n,y,b)^2] = \frac{1}{2^n} \sum_{x \in \{0,1\}^n} \Delta_l(x,y,b)^2 = \frac{4}{2^n} \sum_{j \in \{0,1\}^{u_l}} \sum_{x \in \{0,1\}^n} |g^{(l)}_{x \circ y \circ b}(1^{t_l} \circ j)|^2,$$

which is the desired relationship.

Then, we show that, for any $1 \leq l \leq k$, $i \in \{0,1\}^{t_l}$, and $j \in \{0,1\}^{u_l}$,

$$\sum_{x \in \{0,1\}^n} \sum_{y \in \{0,1\}} \sum_{b \in \{0,1\}^{p+q}} |g^{(l)}_{x \circ y \circ b}(i \circ j)|^2 = 1.$$

This can also be shown by the following simple calculation. For any $1 \leq l \leq k$, the unitary operation V_l can be represented as

$$\begin{aligned} V_l &= \sum_{x \in \{0,1\}^n} \sum_{y \in \{0,1\}} \sum_{b \in \{0,1\}^{p+q}} V_l |x \circ y \circ b\rangle\langle x \circ y \circ b| \\ &= \sum_{x \in \{0,1\}^n} \sum_{y \in \{0,1\}} \sum_{b \in \{0,1\}^{p+q}} \sum_{i \in \{0,1\}^{t_l}} \sum_{j \in \{0,1\}^{u_l}} g^{(l)}_{x \circ y \circ b}(i \circ j) |i \circ j\rangle\langle x \circ y \circ b|. \end{aligned}$$

Thus, the unitary operation $V_l^\dagger$, which is the inverse of V_l, can be represented as

$$V_l^\dagger = \sum_{x \in \{0,1\}^n} \sum_{y \in \{0,1\}} \sum_{b \in \{0,1\}^{p+q}} \sum_{i \in \{0,1\}^{t_l}} \sum_{j \in \{0,1\}^{u_l}} g^{(l)}_{x \circ y \circ b}(i \circ j)^* |x \circ y \circ b\rangle\langle i \circ j|.$$

This implies that, for any $i \in \{0,1\}^{t_l}$ and $j \in \{0,1\}^{u_l}$,

$$V_l^\dagger |i \circ j\rangle = \sum_{x \in \{0,1\}^n} \sum_{y \in \{0,1\}} \sum_{b \in \{0,1\}^{p+q}} g^{(l)}_{x \circ y \circ b}(i \circ j)^* |x \circ y \circ b\rangle.$$

A direct calculation shows that

$$1 = \langle i \circ j | V_l V_l^\dagger | i \circ j \rangle = \sum_{x \in \{0,1\}^n} \sum_{y \in \{0,1\}} \sum_{b \in \{0,1\}^{p+q}} |g^{(l)}_{x \circ y \circ b}(i \circ j)|^2,$$

which is the desired relationship.

The above two relationships imply Lemma 6 as follows. For any $y \in \{0,1\}$ and $b \in \{0,1\}^{p+q}$,

$$\begin{aligned} \sum_{l=1}^k \mathrm{E}[\Delta_l(U_n,y,b)^2] &= \frac{4}{2^n} \sum_{l=1}^k \sum_{j \in \{0,1\}^{u_l}} \sum_{x \in \{0,1\}^n} |g^{(l)}_{x \circ y \circ b}(1^{t_l} \circ j)|^2 \\ &\leq \frac{4}{2^n} \sum_{l=1}^k \sum_{y \in \{0,1\}} \sum_{b \in \{0,1\}^{p+q}} \sum_{j \in \{0,1\}^{u_l}} \sum_{x \in \{0,1\}^n} |g^{(l)}_{x \circ y \circ b}(1^{t_l} \circ j)|^2 \\ &= \frac{4}{2^n} \sum_{l=1}^k \sum_{j \in \{0,1\}^{u_l}} 1 = \frac{4}{2^n} \sum_{l=1}^k 2^{u_l} = \sum_{l=1}^k \frac{2^{p+q+3}}{2^{t_l}} \leq \frac{k 2^{p+q+3}}{2^{t_{\min}}}, \end{aligned}$$

which is the desired first relationship. In particular, it holds that

$$\frac{4}{2^n} \sum_{l=1}^k \sum_{b \in \{0,1\}^{p+q}} \sum_{j \in \{0,1\}^{u_l}} \sum_{x \in \{0,1\}^n} |g^{(l)}_{x \circ 0 \circ b}(1^{t_l} \circ j)|^2 \leq \frac{k 2^{p+q+3}}{2^{t_{\min}}}.$$

We consider the value

$$\sum_{l=1}^{k} \mathrm{E}[\Delta_l(U_n, 0, 0^p \circ a)^2] = \frac{4}{2^n} \sum_{l=1}^{k} \sum_{j \in \{0,1\}^{u_l}} \sum_{x \in \{0,1\}^n} |g^{(l)}_{x \circ 0 \circ 0^p \circ a}(1^{t_l} \circ j)|^2 \tag{6}$$

for any $a \in \{0,1\}^q$. There exists some $a' \in \{0,1\}^q$ such that it minimizes this value, i.e., value (6) with a' is less than or equal to that with any other $a \in \{0,1\}^q$. It holds that

$$2^q \sum_{l=1}^{k} \mathrm{E}[\Delta_l(U_n, 0, 0^p \circ a')^2] \leq \sum_{a \in \{0,1\}^q} \sum_{l=1}^{k} \mathrm{E}[\Delta_l(U_n, 0, 0^p \circ a)^2] \leq \frac{k 2^{p+q+3}}{2^{t_{\min}}},$$

which immediately implies the desired second relationship. □

A.7 Proof of Theorem 3

Proof. We assume on the contrary that there exists a depth-d quantum circuit C_n for PA_n with n input qubits, one output qubit, $p = O(\log n)$ initialized ancillary qubits, and $q = n^{O(1)}$ uninitialized ancillary qubits such that it consists of one-qubit gates, CNOT gates, unbounded fan-out gates on $(\log n)^{O(1)}$ qubits, and unbounded Z gates, where $d = O((\log n)^\delta)$ for some constant $0 \leq \delta < 1$. Since $p = O(\log n)$, there exists a constant $c > 0$ such that $p \leq c \log n$ when n is sufficiently large. We define $t = (c+4)\log(n+p+q+1)$ and consider the set $\mathcal{G}_t$ of all unbounded Z gates in C_n that act on more than or equal to t qubits.

We first consider the case where $\mathcal{G}_t = \emptyset$. Since $t = O(\log n)$ and unbounded fan-out gates act on $(\log n)^{O(1)}$ qubits, all gates in C_n act on at most $w = (\log n)^{O(1)}$ qubits. Therefore, by Lemma 4, the output of C_n can depend only on the states of at most w^d input qubits. Since $d = O((\log n)^\delta)$, it holds that $w^d = o(n)$. Thus, there exists an input qubit of C_n such that the output of C_n does not depend on the input qubit. This implies that C_n cannot compute PA_n since the output of PA_n changes if any one of the n input bits changes. This is a contradiction.

The remaining case is when $\mathcal{G}_t \neq \emptyset$. In this case, we can apply the analysis of a general circuit in Section 4.2. Recall that $\mathcal{G}_t = \{T_1, \ldots, T_k\}$, t_l is the number of qubits on which T_l acts, and $t_{\min} = \min\{t_l | 1 \leq l \leq k\}$. It holds that $p \leq c \log n$, $k \leq (n+p+q+1)d/t_{\min}$, and $t_{\min} \geq (c+4)\log(n+p+q+1)$. Thus, by the second inequality in Lemma 7 with $\varepsilon = 0.1$,

$$\begin{aligned}
\Pr[\Delta(U_n, 0, 0^p \circ a) < 0.1] &\geq 1 - \left(\frac{(n+p+q+1)d}{t_{\min}}\right)^3 \frac{800 n^c}{2^{t_{\min}}} \\
&\geq 1 - \left(\frac{d}{(c+4)\log(n+p+q+1)}\right)^3 \frac{800 n^c}{(n+p+q+1)^{c+1}}
\end{aligned}$$

for some $a \in \{0,1\}^q$. Let us express this value on the right-hand side by $1 - \gamma$. Thus, there exists a set $S \subseteq \{0,1\}^n$ such that S has at least $2^n(1-\gamma)$ elements and, for any $x \in S$, it holds that $\Delta(x, 0, 0^p \circ a) < 0.1$. Since γ goes to 0 as n goes to infinity, it holds that $2^n(1-\gamma) > 2^{n-1}$ when n is sufficiently large. A simple calculation shows that, for any $x \in \{0,1\}^n$ satisfying $\Delta(x, 0, 0^p \circ a) < 0.1$, the output of $\tilde{C}_n|x \circ 0 \circ 0^p \circ a\rangle$ coincides with that of $C_n|x \circ 0 \circ 0^p \circ a\rangle$ with probability of at least $1 - 0.1^2 = 0.99$ [3, 2]. Moreover, when the initial state of the uninitialized ancillary qubits is set to $|a\rangle$, C_n computes PA_n, i.e., for any $x \in \{0,1\}^n$, the output of $C_n|x \circ 0 \circ 0^p \circ a\rangle$ is $\mathrm{PA}_n(x)$. Thus, for any $x \in S$, the output of $\tilde{C}_n|x \circ 0 \circ 0^p \circ a\rangle$ is $\mathrm{PA}_n(x)$ with probability of at least 0.99. This contradicts the fact obtained by the following argument. Since all gates in $\tilde{C}_n$ act on at most $(\log n)^{O(1)}$ qubits, as described for the case where $\mathcal{G}_t = \emptyset$, by Lemma 4, there exists an input qubit of $\tilde{C}_n$ such that the output of $\tilde{C}_n$ does not depend on the input qubit. This implies that, for at most 2^{n-1} elements $x \in \{0,1\}^n$, the output of $\tilde{C}_n|x \circ 0 \circ 0^p \circ a\rangle$ is $\mathrm{PA}_n(x)$ with probability greater than 0.5. □